

Bi³: A Biplatform, Bicultural, Biperson Dataset for Social Robot Navigation*

Andrew Stratton¹, Phani Teja Singamaneni^{2,3}, Pranav Goyal¹, Rachid Alami², and Christoforos Mavrogiannis¹

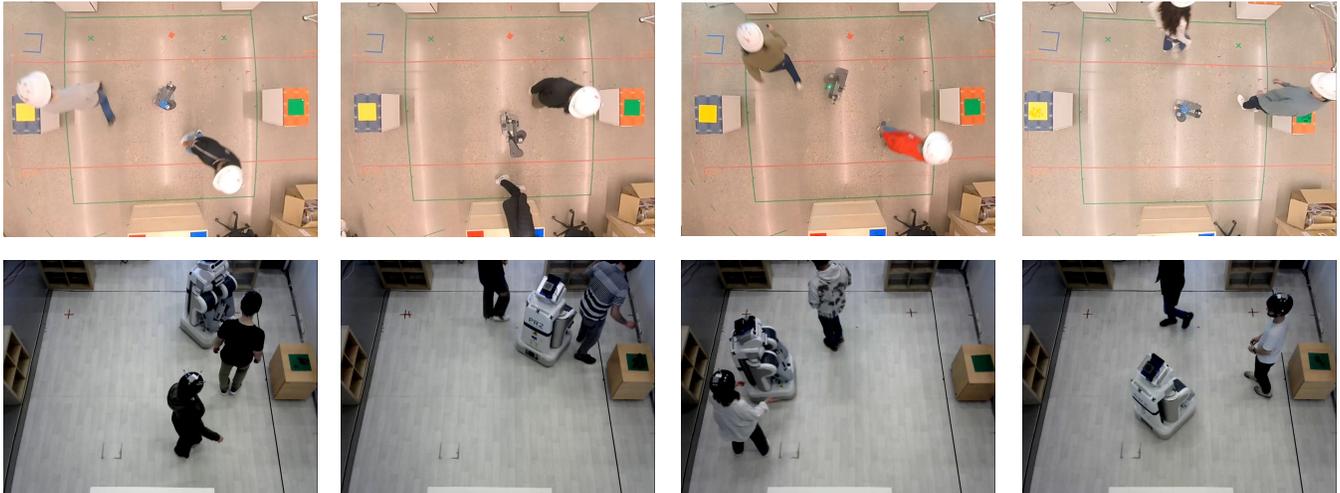


Fig. 1: We introduce Bi³, a novel dataset which contains diverse navigation interactions in close-quarters settings captured with two robots at two sites with five fully autonomous robot controllers. The top row shows example user interactions with a Hello Robot Stretch at UM, whereas the bottom row shows interactions with a Willow Garage PR2 at LAAS-CNRS.

Abstract—We contribute Bi³, a dataset of social robot navigation among groups of people in a constrained lab space. Compared to prior data collection efforts for social robot navigation, our dataset is unique in that it features: an original experiment design giving rise to close navigation encounters between two humans and a robot; five different navigation algorithms; two different robot platforms; a diverse participant pool of 74 people recruited from two sites in the USA and France; multimodal data streams including 10.5 hours of human and robot ground-truth motion tracks, RGB video, and user impressions over robot performance. Our analysis of the collected dataset through metrics like interaction density and human velocity suggests that Bi³ represents a benchmark of unique diversity and modeling complexity. Bi³ contributes towards understanding how humans and robots can productively mesh their activities in constrained environments, and can be a resource for training models of human motion prediction and robot control policies for navigation in densely crowded spaces.

I. INTRODUCTION

A critical component of any social robot navigation (SRN) system is a model for how people move around a

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

²Phani Teja Singamaneni and Rachid Alami are affiliated with LAAS-CNRS, University of Toulouse, Toulouse, FR. {ptsingaman, rachid.alami}@laas.fr.

³Phani Teja Singamaneni is also affiliated with INRIA, University of Lorraine, Nancy, FR. {phani-teja.singamaneni@inria.fr}

*Our dataset is available at: <https://fluentrobotics.com/bi3dataset/>.

This work was partially supported by the Horizon Europe Framework Programme grant 101070596 - euROBIN.

robot [25, 43]. This has motivated extensive research on human motion prediction (HMP) [9, 29, 36–38], the problem of estimating future interactions among groups of people given observations of their behavior over a window of the past. Much of the recent work on HMP makes use of deep-learning techniques due to their ability to represent complex interactions in high-dimensional spaces. The success of these techniques relies on the closeness between the dataset they are trained on and the distribution of interactions they encounter during deployment. This has motivated extensive data collection efforts involving the deployment of cameras, robots, or sensorized devices in crowded human spaces [2, 18, 22, 28, 31, 53, 56].

We observe a few critical gaps in existing datasets supporting SRN and HMP. First, the vast majority of datasets often feature very sparse and distant interactions among people [1, 41]. To truly integrate in realistic human spaces, robots need to be prepared to respond adequately to close human-robot navigation encounters. Second, many datasets lack robot deployments [2, 18, 26, 28, 31, 56], and instead include data collected via overhead cameras. Robots deployed in the field, however, will need to leverage onboard sensing to perceive and react to people, leading to additional perception challenges [22, 56]. When robots are part of a data collection effort, they are often deployed “in the wild.” While it is important to train systems on real-world, unscripted interactions, the lack of control for external variables influencing human behavior around robots could lead to

TABLE I: Comparison between our dataset and other HMP and SRN datasets. Blank columns are not applicable to that dataset row.

| Dataset | Scene | Multiple Sites | Duration | Pose Tracking | Goals | Multiple Robot Behaviors | # of Robot Platforms | User Impressions |
|------------------------|----------------|----------------|----------|---------------|-------|--------------------------|----------------------|------------------|
| ETH [28] | Street | ✓ | 25' | Camera | | | 0 | |
| UCY [18] | Street | | 20' | Camera | | | 0 | |
| Edinburgh [20] | Outdoor | | 4 months | Camera | | | 0 | |
| ATC [2] | Indoor | | 41 days | 3D range | | | 0 | |
| SDD [31] | Outdoor | ✓ | 5 days | Camera | | | 0 | |
| TBD [53] | Indoor | | 12.7 hrs | Camera | | | 1 | |
| JRDB [22] | Indoor/outdoor | ✓ | 64' | Camera | | ✓ | 1 | |
| L-CAS [56] | Indoor | ✓ | 49' | 3D LiDAR | | | 1 | |
| SCAND [14] | Indoor/outdoor | ✓ | 8.7 hrs | | | | 2 | |
| MuSoHu [26] | Indoor/outdoor | ✓ | 20 hrs | | | | 0 | |
| INTERACT [57] | Lab | | 9 hrs | MoCap | ✓ | | 1 | |
| THOR [32] | Lab | | 1 hr | MoCap | ✓ | | 1 | |
| THOR MAGNI [40] | Lab | | 3.5 hrs | MoCap | ✓ | ✓ | 1 | |
| HuRON [13] | Office | | 75 hrs | | | ✓ | 1 | |
| CrowdBot [27] | Outdoor | | 1 hr | | | ✓ | 1 | |
| Bi ³ (Ours) | Lab | ✓ | 10.5 hrs | MoCap | ✓ | ✓ | 2 | ✓ |

poor data associations performed during model training [7]. Many robot-deployment datasets are collected via remote teleoperation of a robot, without any autonomy involved [14, 22, 40]. This does not allow for the extraction of insights about how algorithm embodiment would impact human-robot interaction. Many datasets recording human-robot navigation encounters lack any information about the impressions of people around them. This impedes the extraction of insights about how to assess and design *socially competent* robot behavior, which is ultimately the main goal of SRN [25]. Finally, the perceived social competence of robot behavior is often a factor of robot morphology and cultural norms present in a deployment region [52]. Thus, it is important to enable researchers to introduce these dependencies into the model training and validation.

Motivated by these gaps (see table I for a summary), we introduce Bi³, a new SRN dataset collected during our IRB-approved user study [45], featuring: two different platforms available to numerous labs around the world (Hello Robot Stretch, Willow Garage PR2, see fig. 1); 10.5 hours of natural, close navigation encounters between a robot and 74 participants, recruited from two sites in distinct regions of the world (University of Michigan in Ann Arbor, USA and LAAS-CNRS in Toulouse, France); a wide range of autonomous robot behavior resulting from the embodiment of five different HMP models on a model predictive control (MPC) scheme; high-accuracy human and robot motion tracks; top-down RGB video and egocentric depth images; prediction-control pairs for every control cycle; user impressions about the robot collected via Likert-style questionnaires and open-ended comments. Our analysis shows that Bi³ has very dense, high-velocity human interactions, diverse human and robot motion, and diverse human preferences.

II. RELATED WORK

Initial dataset efforts for analyzing human navigation interactions, including the smaller-scale ETH [28] and UCY [18] datasets, as well as the larger-scale Edinburgh [20], ATC [2] and Stanford Drone (SDD) [31] datasets, focused mainly on open, public spaces. While these datasets do not include robots in their scenes, and are thus less suited to

studying human-robot navigation interactions, they have been used extensively for the training and benchmarking of HMP algorithms [1, 9, 36–38] and multi-target tracking [6, 35].

Several datasets have been collected in similar spaces using robots or sensorized devices. The TBD dataset [53] contributes to the data collection pipeline with a portable data collection and automated labeling system. The JRDB [22] and L-CAS [56] datasets collected sensor data aboard a mobile robot to enable studying robot-centric perception challenges, including tracking [51, 56], activity recognition [5], as well as HMP [34, 38]. The SCAND [14] dataset leveraged teleoperation with two mobile robots to collect expert SRN demonstrations for training imitation learning policies, while the MuSoHu [26] dataset has human demonstrators equipped with sensors. These datasets, owing to their collection in public areas with ample space, are sparse in regard to close-quarters human-robot interactions [1]. Bi³ is collected in a constrained workspace with structured human and robot tasks that forced consistent close encounters, making it more effective for learning models for deployment in many indoor settings where space is limited.

The datasets most similar to ours are those which are collected through experiments in lab settings, allowing for finer control over the interactions occurring in the data. The INTERACT [57] dataset has three humans interacting with a robot, but focuses on human skeleton trajectory prediction during collaborative assembly tasks. Our work instead uses assembly tasks to motivate challenging multi-agent navigation interactions. The THOR [32] and THOR MAGNI [40] datasets also focus on multiagent interactions resulting from humans and a robot moving between stations to complete tasks in a shared workspace. They prioritize high task diversity, larger numbers of humans, and many sensor modalities in support of HMP with additional human features [19, 33, 58]. Our dataset instead adds diversity of close interactions by recording human-robot navigation encounters resulting from five different algorithms, in two locations, and with two robot platforms, factors which increase the applicability of our data to a wider range of downstream robot deployments. Crucially, our dataset also includes human-centered attributes through the recording of

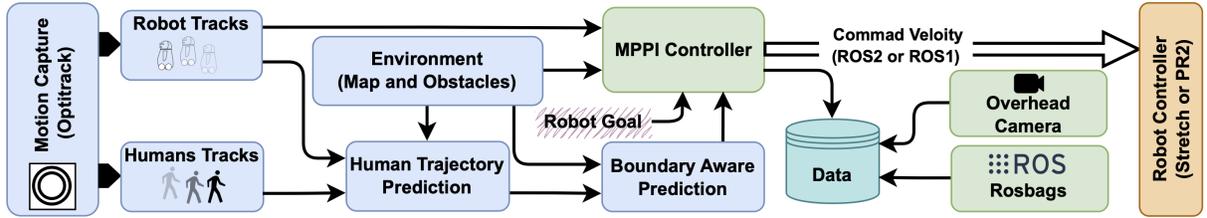


Fig. 2: Architecture and data collection pipeline.

user impressions of robot sociability [3] and workload [12] which can enable the development and validation of human-centered algorithms.

III. SOCIAL ROBOT NAVIGATION SYSTEM

Our SRN system is implemented as an MPC architecture due to its flexibility, modularity, and community adoption [10, 29, 39, 44]. We formalize our domain, and discuss the HMP models used in our data collection. A diagram of our system and data collection pipeline is in figure 2.

A. Domain

We consider a robot navigating among $n \geq 1$ humans in a workspace $\mathcal{W} \subseteq \mathbb{R}^2$. We describe the state of the robot as $s^r \in \mathcal{W}$ and the state of human $i \in \{1, \dots, n\}$ as s^i . The state of the robot evolves according to dynamics $s_{t+1}^r = g(s_t^r, u_t)$, where u_t is a control action (speed, and steering angle), drawn from a space of controls \mathcal{U} . The robot starts from an initial configuration s_0^r and moves towards a goal d^r by following a policy π^r , while humans navigate from their initial configurations s_0^i towards their goals, d^i by following a policy π^i , $i \in \mathcal{N}$; agents' goals and policies are unknown to one another. We assume that the robot perfectly observes human states at all times. We focus on the problem of designing a robot policy π^r , such that the robot reaches its goal d^r in a safe, efficient, and socially compliant fashion.

B. Model Predictive Control for Social Robot Navigation

Our MPC is an implementation of the robot policy π^r as a receding-horizon control optimization over a horizon T :

$$\begin{aligned}
 u_{t:T}^* &= \arg \min_{u_t \in \mathcal{U}} \sum_t^{T-1} \mathcal{J}(s_{t+1}^r, s_{t+1}^{1:n}) \\
 \text{s.t. } s_{t+1}^r &= g(s_t^r, u_t) \\
 s_{t+1}^{1:n} &= f(s_{t-h:t}^r, s_{t-h:t}^{1:n}),
 \end{aligned} \tag{1}$$

where \mathcal{J} is a cost function capturing qualities of SRN, and f is an HMP model that takes as input the state history of the robot $s_{t-h:t}^r$ and humans $s_{t-h:t}^{1:n}$ over a window h , and outputs a prediction of human motion, $s_{t+1}^{1:n}$. Following prior work [24, 29], we model cost \mathcal{J} as:

$$\mathcal{J}(s^r, s^{1:n}) = a_g \mathcal{J}_g(s^r) + a_d \mathcal{J}_d(s^r, s^{1:n}) + a_o \mathcal{J}_o(s^r), \tag{2}$$

where $\mathcal{J}_g(s^r) = (s^r - d^r)^2$ penalizes distance to the robot's goal, $\mathcal{J}_d(s^r, s^{1:n}) = \sum_{i=1}^n A_d^2(s^r, s^i)$, penalizes violations to the personal space as done in prior work [15], and $\mathcal{J}_o(s^r) = \mathbb{1}(A_R(s^r) \cap \mathcal{W}^c \neq \emptyset)$ penalizes violations of the workspace boundary, where $A^r(s^r)$ is the volume occupied by the robot at pose s^r . Weights a_g , a_d , and a_o represent the relative importance of each cost.

C. Human Modeling

Our system's modular design allows the integration of any HMP model. To capture a wide diversity of robot human-robot navigation encounters (examples are shown in fig. 3), we integrated five different HMP models into our MPC:

No Prediction (NP): The robot does not perceive the humans navigating around it. NP is representative of the common practice in industrial robotics of having unreactive robots [11].

Static (ST): Humans are treated as static obstacles, as in many costmap-based SRN approaches [30, 48]. In close-quarters scenarios in which agents rely on each other moving out of their paths to make progress, ST represents a highly conservative model.

Constant Velocity (CV): Humans are treated as dynamic obstacles with no human-aware modeling. CV has repeatedly been shown to be an effective model of human navigation not only in terms of prediction accuracy [41], but also for downstream use in SRN [16, 29, 44, 48].

Human Scene Transformer (HST): HST [38] is a transformer-based generative model which represents the large class of deep learning-based prediction methods [9, 36–38]. We choose the HST for its strong performance on established benchmarks [18, 22, 28].

CoHAN: CoHAN [42, 47, 48] is a trajectory planner that plans robot and predicted human trajectories jointly, representing the paradigm of *cooperative collision avoidance* [23, 42, 46–50]. It optimizes the trajectories using social costs including human-human and human-robot safety, human visibility and speed, and obstacle avoidance.

D. Implementation Details

Hardware: We instantiated our MPC under all HMP models on two different robot platforms: a Hello Robot Stretch at UM, and a Willow Garage PR2 at LAAS. The two platforms are opposite in their size ($330 \times 340 \times 1410mm$ v.s. $668 \times 668 \times 1330mm$) and mass (24.5 kg v.s. 226.8 kg), which we anticipated would affect human perceptions and navigation interactions. Overhead motion capture systems (OptiTrack at both sites) provided high-accuracy robot localization and people tracking at 120 Hz. The coverage spanned the entire workspace at both sites; thus, the commonly reported *Tracking Duration* metric [2, 18, 28, 32, 40, 53] is simply the average length of our trials, which are 183.67s at UM and 226.22s at LAAS. HMP models and MPC were verified and tuned on hardware at each site to ensure consistent behavior. Fig. 5 shows the equipment used at both sites to collect the data.

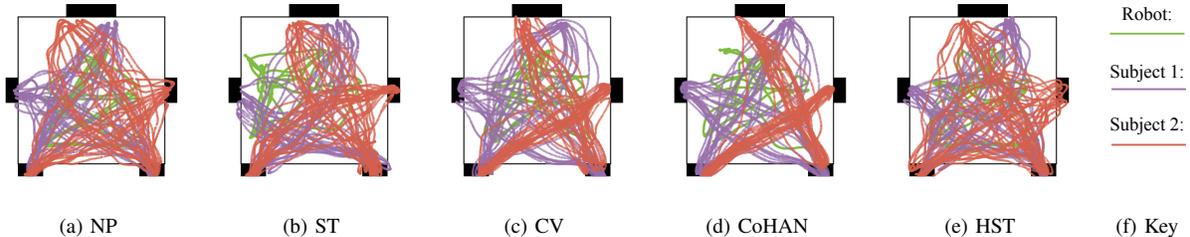


Fig. 3: Human and robot paths over a single trial for all five controllers. Despite the humans and robot having the same sequence of workspace goals for all five, the resulting paths are substantially different.

HMP: All models used an $h = 8$ timestep history and a prediction horizon of $T = 12$, with a timestep of $0.4s$, matching standard evaluation settings in the HMP literature [9, 36–38]. Predictions violating the workspace boundary were truncated and replaced with the latest valid within-boundary predictions. The HST checkpoint was trained on the ZARA1 portion of the ETH/UCY dataset, producing $m = 20$ prediction modes, the most likely of which was used. To match other models, CoHAN’s backoff behavior and FoV restriction on planning were removed.

MPC: Our control architecture is based on a GPU-parallelized implementation of Model Predictive Path Integral (MPPI) control [55], using 2000 trajectory samples. The maximum speed of the robot was capped to $0.3m/s$, as it was found to be a safe limit for our setup. The controller uses a differential-drive kinematic model with a control space consisting of linear and angular velocities at both sites. The full prediction and control loop closed at different frequencies depending on the computational burden of the HMP model, with NP, ST, and CV operating at 50 Hz, HST at 20 Hz, and CoHAN at 8 Hz.

Computation and communications: All computation was performed on lab computers, with velocity commands and sensor data transmitted wirelessly. All communication was done using ROS (ROS2 for Stretch and ROS1 for PR2) with the MPPI running in ROS2 and CoHAN in ROS1. This required ROS bridges to convert messages between the two ROS versions for the two robot platforms.

Software: The system code is modularized so prediction component can be easily swapped with any trajectory prediction model operating on positions. Furthermore, the controller implementation is not platform specific: it works with any low-level controller which executes velocity commands.

IV. DATA COLLECTION

We conducted an ethics approved user study (University of Michigan IRB HUM00259961; Université Toulouse Ethics Approval 2024_917) to collect our data through controlled experiments. For the full details and in-depth analysis of human-robot interactions and user impressions, we refer the reader to our user-study paper [45]. Here, we summarize our scenario design and detail our data collection procedure.

A. Scenario Design

Collecting data on group interactions with a navigating robot in a lab environment is challenging: compared to data collections in the field, more care must be taken to ensure

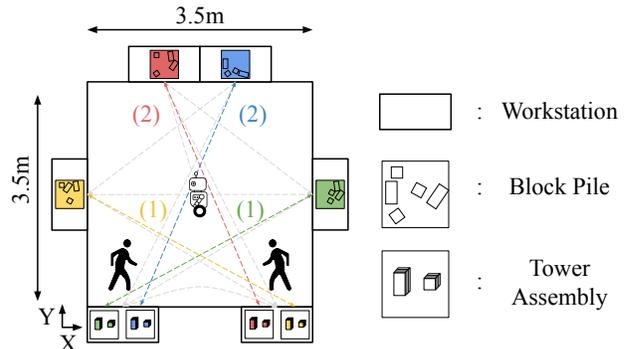


Fig. 4: Experimental setup. Subjects begin and end in front of the assembly stations while the robot begins in the center of the workspace. The blocks labeled (1) are assembled by subjects followed by those labeled (2), while the robot traverses between the goals along the path in gray.

natural human motion and diverse interactions even during highly structured scenarios. To that end, we build upon prior work instantiating industrial teamwork settings in a lab environment [17, 23].

Our scenario involves two subjects repeatedly traversing an open workspace between workstations along the perimeter, picking blocks from piles, traversing, and placing them one-by-one to assemble towers. To match many industrial and home scenarios which necessitate close-quarters interactions, we made the traversable portion of the workspace $3.5 \times 3.5m$. To limit subject interactions to navigation, participants were each responsible for their own set of two out of four total towers and instructed not to talk. These sets, along with a specified ordering for completing the towers in each set, were designed to maximize interaction density by ensuring participants consistently crossed paths with each other and the robot as shown in Fig. 4.

While subjects assembled towers, the robot navigated through a randomly generated sequence of workstations. To instigate regular crossings with both subjects, the sequence required the robot to cross the y-axis of the workspace between each goal (Fig. 4). The sequence was flipped across the y-axis each trial to reduce learning effects and further increase the diversity of interactions. To give participants the impression the robot also completed tasks even though it did not physically interact with the environment, they were told the robot was monitoring their progress.

B. Collection Procedure

In each experiment, two subjects completed five trials, each with a distinct HMP model. Subjects began by reading



Fig. 5: Hardware used for the data collection.

and signing a consent form, after which they were briefed on the scenario and completed a practice trial in which the robot remained static at its starting position. They then completed the five trials, with impressions being collected via online surveys in between each trial. To account for ordering effects, the condition order was determined using a balanced Latin square. Upon completing all five trials and surveys, subjects completed demographic surveys, were given compensation or gifts, and debriefed with an information sheet describing the specific details of the experiment.

Subjects started and ended each trial at ‘X’ shapes marked on the ground in front of the two assembly stations, while the robot began in the center of the workspace facing towards the subjects. The task began once verbal confirmation of readiness was received from both subjects, and the robot began moving. It was considered complete when both subjects completed their towers and returned to their starting positions. Each experiment took approximately 50 minutes and yielded approximately 17 minutes of data.

V. THE Bi^3 DATASET

The dataset consists of 37 experiments (17 taking place at UM and 20 at LAAS) involving a total of 74 subjects. Sample paths from the dataset are shown in Fig. 3. Bi^3 can be downloaded at: <https://fluentrobotics.com/bi3dataset/>.

A. Objective Data

Our dataset includes raw sensor data including robot and human transforms from motion-capture as well as video recorded from cameras at both sites (Insta360 at UM and Logitech C270 at LAAS). Both cameras faced top-down, with the UM camera in the center of the workspace and the LAAS camera at one side. Video was synchronized with saved ROS bags (ROS2 in UM and ROS1 in LAAS) using recording metadata and message timestamps. Data from UM also includes depth camera data from the Stretch’s onboard Intel Realsense D435i, which is mounted at the top of the robot. Top-down camera data was captured at 30 Hz at both sites, while the egocentric camera published at 15 Hz. In addition to raw sensor data, we captured additional information calculated as part of the control code, including HMP model outputs, the robot goal, and robot velocity

commands. These were stored in JSON files, saved directly at each control loop during the experiment. Finally, we contribute auxiliary information regarding the locations of the workspace boundary, workstations, human goals, and human tasks, which have been shown to improve prediction performance in similar domains [40].

B. User Impressions

We move beyond prior SRN datasets by also including human impressions of the robot performance per HMP model. Human impressions have emerged as not only critical for human-robot interaction studies, but also as a modality for training data-driven models for interactive tasks [4, 54]. Our surveys included quantitative impressions, collected via the Robotic Social Attributes Scale (RoSAS) [3], and the NASA Task Load Index (TLX) [12]. The full RoSAS survey consists of 18 Likert-style questions which rate the amount an attribute applied to the robot. Six questions capture *Discomfort*, six capture *Competence*, and six capture *Warmth*; we omitted the Warmth questions because they do not readily apply to SRN interactions, and used a nine-point scale. The TLX consists of 21-point Likert-style questions in which the user must rate their perceived workload along six different axes; we used all six questions. Our surveys also collected qualitative impressions via short open-ended questions requesting optional additional thoughts. These have been shown to enable the extraction of critical insights about user experience that are not well captured by standard scales [23]. The dataset contains 285 survey responses and 78 comments from 57 participants (17 at UM, 40 at LAAS), as participants at UM were able to opt out of including their survey responses.

VI. ANALYSIS

We discuss the quality and diversity of the human motion, robot motion, and human impressions of Bi^3 .

Human Motion Profile. Our data collection, through its large bicultural subject pool, diverse robot behaviors, and distinct robot embodiments, generated a highly diverse set of human navigation behaviors. We quantify this through the following metrics proposed in prior work [32]:

- *Minimum Distance (m)*: The minimum distance between the two humans over a single tracklet, which is commonly used in SRN work related to personal space, or *proxemics* [8, 23, 44]. Lower minimum distance indicates closer interactions.
- *Velocity (m/s)*: The velocity of each human. Lower velocities allow more time for other agents to adapt their trajectories cooperatively, allowing for potentially smoother interactions.

We compare against related datasets in figure 6. We verify the high diversity of human motion in Bi^3 through its substantially higher standard deviation of recorded human velocity. Additionally, our data has the lowest minimum distance at both sites and the highest velocity at LAAS. These attributes make human trajectories difficult to predict, demonstrating the challenge Bi^3 presents.

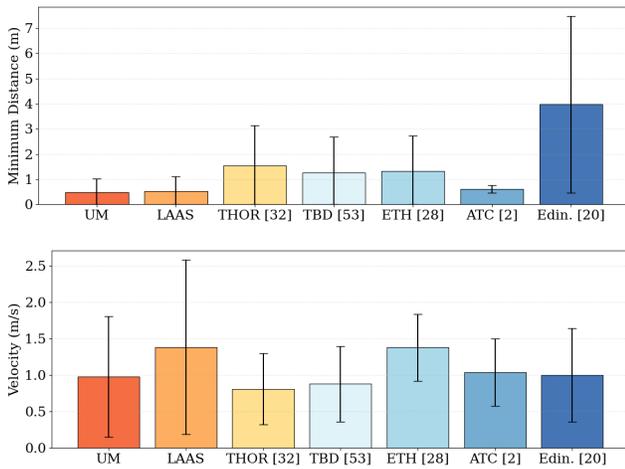


Fig. 6: Human motion statistics between our dataset and several others (bars indicate means and errors are standard deviations).

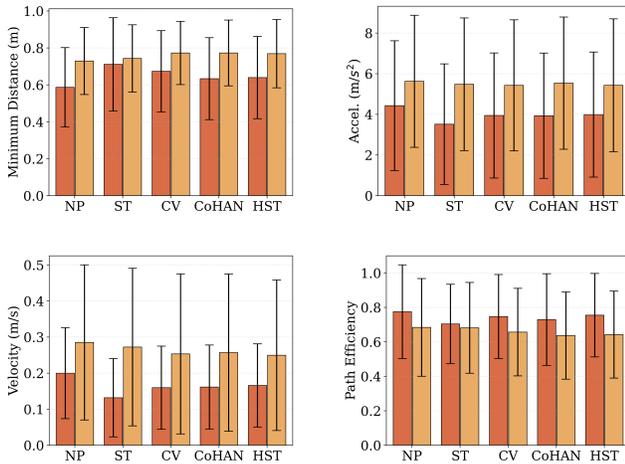


Fig. 7: Robot motion statistics by controller and site (bars are means and errors are standard deviations). Dark orange refers to UM and light orange to LAAS.

Robot Motion Profile. We quantify the differences in robot motion arising from distinct HMP models using the following metrics common to the SRN literature [8, 23, 25]:

- *Minimum Distance (m)*: Higher Minimum Distance is often connected to discomfort or unsafe behavior in prior work [8].
- *Acceleration (m/s²)*: Higher average acceleration is indicative of more sudden movements, which have the potential to disrupt or startle co-navigating humans.
- *Velocity (m/s)*: Higher velocities indicate the robot moved efficiently, with less time slowed for humans.
- *Path Efficiency (ratio)*: The ratio of the distance between the trajectory start and endpoints and the cumulative distance traveled. Higher Path Efficiency can indicate less cooperative robot motion, as the robot does not adjust its paths for humans.

Averages over the above robot metrics by algorithm and by site are in figure 7 (dataset names correspond to Fig. I). We identify that the robot had consistently higher distances from humans, faster and less smooth movement, and less direct paths at LAAS. Our control over the experimental setup leads us to speculate that the larger platform at LAAS caused the

robot to need longer paths to avoid humans, or that cultural differences between the two sites may have played a role.

Within each site we see more limited diversity between controllers. In particular, we find that NP and ST have distinctly higher and lower values across all metrics than CV, CoHAN, and HST. Despite the differences in motion metrics for CV, CoHAN, and HST being more subtle, we find that they received vastly different user impressions. This highlights the importance of user impressions for identifying desirable SRN algorithm behavior.

User Impressions and Comments. We see that each robot behavior elicited vastly different impressions from users. The RoSAS and NASA TLX survey responses can be found in tables II and III respectively. We see substantial differences in mean both between algorithms and between sites, indicating each controller was received very differently by users. Furthermore, the consistently high standard deviations by trial serve to indicate that impressions varied strongly between each individual trial, a fact which shows the diversity of impressions exhibited in our dataset.

User comments also provide valuable descriptions of specific behavior that informed users’ impressions of the robot. High level themes are substantially different by model. NP and CoHAN received the most negative comments, with many concerning them being “*less responsive*” and “*too close*.” ST and CV received more balanced comments, however they still received comments describing them as “*hard to predict*” (ST) and that “*the robot ended up staying in annoying places*” (CV). HST received the most positive and least negative comments, with specific comments describing it as “*very interactive*” and that it “*was a lot smoother compared to the other ones*.” Notably, all models received at least some positive and negative comments, which can help identify scenarios in which models that performed worse overall still achieved desirable behavior.

VII. ANTICIPATED USE CASES

Human motion prediction in close-quarters human-robot interactions: HMP in close-quarters, task-driven settings is particularly challenging due to the highly irregular paths caused by forced close interactions, high velocities caused by the need for productivity, and rapid direction switches caused by moving between specific goals. Deployment in warehouses, hospitals, and airports makes high performance under such settings necessary. As demonstrated in Sec. VI, our dataset captures these challenges, and thus presents a quality setting for developing such HMP algorithms. Furthermore, the human and robot goal information we include can aid in developing goal conditioned prediction models [21] which may prove better in these settings.

Learning robust social navigation policies. The inclusion of a robot navigating among humans makes this dataset valuable for directly learning SRN policies or models which can be integrated into MPCs. In addition to containing desirable behavior, our dataset contains information on how humans react when the robot acts in error—valuable data for learning policies that are robust to failure. Methods such as

TABLE II: RoSAS scale responses by algorithm (mean and standard deviation). The top row for each attribute refers to UM and the bottom row to LAAS. **Discomfort** is a mean of the first six rows, whereas **Competence** is a mean of the bottom six.

| Attribute | NP | ST | CV | CoHAN | HST |
|-------------------|-----------|-----------|-----------|-----------|-----------|
| Aggressive | 2.53±2.15 | 2.06±1.43 | 2.59±1.84 | 2.88±2.42 | 2.59±2.15 |
| | 2.80±2.20 | 1.85±1.46 | 1.90±1.43 | 3.22±2.42 | 2.05±1.57 |
| Awful | 1.71±1.40 | 1.82±1.51 | 1.47±0.80 | 1.82±1.33 | 1.71±1.05 |
| | 2.72±1.89 | 2.48±1.99 | 2.22±1.78 | 2.35±1.79 | 1.92±1.47 |
| Awkward | 4.65±1.87 | 5.06±2.33 | 4.76±2.33 | 4.88±2.09 | 5.12±2.50 |
| | 4.00±2.12 | 2.82±2.11 | 3.05±2.10 | 3.72±2.17 | 3.25±2.06 |
| Dangerous | 2.06±1.68 | 1.65±1.00 | 2.00±1.66 | 2.24±1.82 | 2.35±1.97 |
| | 2.48±1.99 | 2.00±1.66 | 1.82±1.26 | 3.25±2.17 | 1.85±1.33 |
| Scary | 1.82±1.55 | 1.71±1.10 | 1.82±1.42 | 1.94±1.71 | 1.88±1.62 |
| | 2.72±2.04 | 2.02±1.66 | 1.88±1.18 | 3.00±2.10 | 1.98±1.46 |
| Strange | 4.00±2.29 | 4.06±2.68 | 4.59±2.37 | 3.94±2.38 | 4.35±2.42 |
| | 4.03±2.24 | 3.32±2.23 | 3.20±1.98 | 3.65±2.05 | 2.60±1.84 |
| Knowledgeable | 4.82±1.51 | 4.65±1.80 | 5.12±1.58 | 4.88±1.69 | 4.53±1.74 |
| | 4.78±1.86 | 4.82±1.71 | 4.90±1.74 | 4.68±1.87 | 5.05±2.02 |
| Interactive | 4.24±2.08 | 4.76±2.02 | 4.24±1.75 | 4.24±2.17 | 3.94±1.78 |
| | 4.30±2.11 | 4.38±2.32 | 4.53±2.20 | 3.62±1.93 | 4.78±2.41 |
| Responsive | 5.12±1.87 | 6.18±2.10 | 4.59±1.62 | 4.94±2.49 | 4.71±2.47 |
| | 5.03±2.04 | 5.70±1.76 | 5.52±2.10 | 4.05±1.62 | 6.12±2.02 |
| Capable | 5.53±1.77 | 5.47±1.42 | 5.47±1.23 | 5.00±1.80 | 5.00±1.87 |
| | 4.95±1.66 | 5.55±1.65 | 5.55±2.09 | 5.00±1.80 | 5.58±1.63 |
| Reliable | 4.94±1.68 | 5.06±1.34 | 5.12±1.45 | 4.76±1.99 | 4.65±1.97 |
| | 5.20±1.87 | 5.38±1.78 | 5.60±1.91 | 4.95±1.83 | 5.38±1.97 |
| Competent | 5.76±1.44 | 4.88±1.73 | 4.82±1.63 | 4.76±1.56 | 4.53±1.77 |
| | 4.80±1.68 | 5.28±1.88 | 5.52±2.00 | 4.60±1.81 | 5.28±2.04 |
| Discomfort | 2.79±1.07 | 2.73±1.06 | 2.87±1.14 | 2.95±1.50 | 3.00±1.38 |
| | 3.12±1.60 | 2.42±1.49 | 2.35±1.05 | 3.20±1.56 | 2.28±1.18 |
| Competence | 5.07±1.32 | 5.17±1.20 | 4.89±0.99 | 4.76±1.50 | 4.56±1.70 |
| | 4.84±1.53 | 5.22±1.36 | 5.27±1.59 | 4.48±1.37 | 5.36±1.46 |

TABLE III: NASA TLX scale responses by algorithm (mean and standard deviation). The top row for each attribute are statistics from UM, while the bottom row are from LAAS.

| Workload Type | NP | ST | CV | CoHAN | HST |
|---------------|-----------|-----------|-----------|-----------|-----------|
| Mental | 2.24±2.91 | 2.12±2.57 | 2.35±3.12 | 2.71±3.04 | 2.82±3.38 |
| | 2.45±2.59 | 2.65±2.57 | 2.40±2.58 | 3.30±3.25 | 2.38±2.70 |
| Physical | 2.82±3.28 | 3.00±2.76 | 3.18±3.73 | 2.94±2.66 | 2.71±2.66 |
| | 2.20±2.37 | 1.92±2.19 | 2.22±2.37 | 2.50±2.59 | 2.08±2.10 |
| Temporal | 3.12±2.64 | 3.94±3.80 | 4.41±3.64 | 3.18±2.30 | 4.41±5.03 |
| | 3.52±3.43 | 2.80±3.60 | 3.15±3.42 | 3.52±3.45 | 2.80±3.22 |
| Performance | 18.6±5.50 | 19.8±1.98 | 18.2±5.32 | 19.5±3.28 | 17.9±5.75 |
| | 18.8±5.08 | 19.2±3.78 | 18.2±5.37 | 18.8±4.18 | 18.9±4.61 |
| Effort | 2.53±2.62 | 3.35±3.66 | 2.59±2.94 | 3.00±3.34 | 2.94±3.54 |
| | 3.32±4.51 | 2.22±2.58 | 2.98±3.17 | 2.90±2.73 | 2.20±2.30 |
| Frustration | 1.94±2.90 | 2.47±3.41 | 2.35±2.40 | 3.24±4.37 | 3.65±5.16 |
| | 3.90±4.33 | 2.30±2.88 | 3.15±3.98 | 4.25±4.64 | 2.38±3.10 |

offline reinforcement learning [7] can be used to learn SRN policies that emulate desirable behavior while simultaneously learning to avoid failure modes in the data.

Learning from human preferences. Existing methods for preference-based learning in SRN have crowdsourced annotations of preferences over trajectory choices online [54]. The lack of physical interaction between labelers and the robot introduces a fundamental discrepancy between the labelers' preferences and those who interact with the robot during deployment. The numerical survey responses in our dataset overcome this by providing impression labels directly from the users in the trajectory data immediately following their real-world SRN interactions. Techniques such as preference-based inverse reinforcement learning [4] can be used with our data to learn reward functions matching the preferences of users who interact with the physical system.

Identifying embodiment and cultural effects in social navigation. It is inevitable that SRN systems will be deployed on different robots in areas with distinct cultural norms. Our dataset contains user interactions with both a lightweight, small robot and a heavy, large robot in two

culturally distinct locations. Although both morphology and culture varying across sites limits our ability to attribute the observed differences to either specifically, the variation over these factors increases the applicability of our data to a broader range of future SRN deployment settings.

REFERENCES

- [1] J. Amirian, B. Zhang, F. V. Castro, J. J. Baldelomar, J.-B. Hayet, and J. Pettré. OpenTraj: Assessing prediction complexity in human trajectories datasets. In *Proceedings of the Asian Conference on Computer Vision (ACCV)*, pages 566–582, 2020.
- [2] D. Bršćić, T. Kanda, T. Ikeda, and T. Miyashita. Person tracking in large public spaces using 3-d range sensors. *IEEE Transactions on Human-Machine Systems*, 43(6):522–534, 2013.
- [3] 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 ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, page 254–262, 2017.
- [4] P. F. Christiano, J. Leike, T. B. Brown, M. Martic, S. Legg, and D. Amodei. Deep reinforcement learning from human preferences. In *Proceedings of the International Conference on Neural Information Processing Systems (NIPS)*, page 4302–4310, 2017.
- [5] M. Ehsanpour, F. Saleh, S. Savarese, I. Reid, and H. Rezatofighi. Jrdback: A large-scale dataset for spatio-temporal action, social group and activity detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [6] F. B. Fuchs, A. R. Kosiorek, L. Sun, O. P. Jones, and I. Posner. End-to-end recurrent multi-object tracking and trajectory prediction with relational reasoning. *ArXiv*, abs/1907.12887, 2019.
- [7] S. Fujimoto, D. Meger, and D. Precup. Off-policy deep reinforcement learning without exploration. In *International Conference on Machine Learning*, 2018.
- [8] Y. Gao and C.-M. Huang. Evaluation of socially-aware robot navigation. *Frontiers in Robotics and AI*, 8, 01 2022.
- [9] 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.
- [10] J. Han, N. Rhinehart, H. Thomas, J. Zhang, and T. D. Barfoot. DR-MPC: Deep residual model predictive control for real-world social navigation. In *IEEE Robotics and Automation Letters*, 2025.
- [11] A. Hanna, S. Larsson, P.-L. Götvall, and K. Bengtsson. Deliberative safety for industrial intelligent human–robot collaboration: Regulatory challenges and solutions for taking the next step towards industry 4.0. *Robotics and Computer-Integrated Manufacturing*, 78:102386, 2022.
- [12] 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.
- [13] N. Hirose, D. Shah, A. Sridhar, and S. Levine. Sacson: Scalable autonomous control for social navigation. *IEEE Robotics and Automation Letters*, PP:1–8, 01 2023.
- [14] H. Karnan, A. Nair, X. Xiao, G. Warnell, S. Pirk, A. Toshev, J. Hart, J. Biswas, and P. Stone. Socially compliant navigation dataset (scand): A large-scale dataset of demonstrations for social navigation. *IEEE Robotics and Automation Letters*, 2022.
- [15] R. Kirby. *Social Robot Navigation*. PhD thesis, Carnegie Mellon University, 2010.
- [16] M. Kollmitz, K. Hsiao, J. Gaa, and W. Burgard. Time dependent planning on a layered social cost map for human-aware robot navigation. In *Proceedings of the European Conference on Mobile Robots (ECMR)*, pages 1–6, 2015.
- [17] L. Landolfi, D. Pasquali, A. Nardelli, J. Bernotat, and F. Rea. Working Memory-Based Architecture for Human-Aware Navigation in Industrial Settings. In *Proceedings of the IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, pages 1878–1885, 2023.
- [18] A. Lerner, Y. Chrysanthou, and D. Lischinski. Crowds by Example. *Computer Graphics Forum*, 2007.
- [19] M. Li, B. Zhong, E. Lobaton, and H. Huang. Fusion of human gaze and machine vision for predicting intended locomotion mode. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 30: 1103–1112, 2022.
- [20] B. Majecka. Statistical models of pedestrian behaviour in the forum. *Master's thesis, School of Informatics, University of Edinburgh*, 2009.

- [21] K. Mangalam, Y. An, H. Girase, and J. Malik. From goals, waypoints & paths to long term human trajectory forecasting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 15213–15222, 2021.
- [22] R. Martin-Martin, M. Patel, H. Rezatofighi, A. Shenoj, J. Gwak, E. Frankel, A. Sadeghian, and S. Savarese. JRDB: A dataset and benchmark of egocentric robot visual perception of humans in built environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- [23] C. Mavrogiannis, P. Alves-Oliveira, W. Thomason, and R. A. Knepper. Social Momentum: Design and Evaluation of a Framework for Socially Competent Robot Navigation. *ACM Transactions on Human-Robot Interaction*, 11(2):14:1–14:37, Feb. 2022.
- [24] 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.
- [25] 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.
- [26] D. M. Nguyen, M. Nazeri, A. Payandeh, A. Datar, and X. Xiao. Toward Human-Like Social Robot Navigation: A Large-Scale, Multi-Modal, Social Human Navigation Dataset. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7442–7447, 2023.
- [27] D. Paez-Granados, Y. He, D. Gonon, D. Jia, B. Leibe, K. Suzuki, and A. Billard. Pedestrian-robot interactions on autonomous crowd navigation: Reactive control methods and evaluation metrics. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 149–156, 2022.
- [28] S. Pellegrini, A. Ess, K. Schindler, and L. Van Gool. You’ll never walk alone: Modeling social behavior for multi-target tracking. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 261–268, 2009.
- [29] 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.
- [30] A. H. Raj, Z. Hu, H. Karnan, R. Chandra, A. Payandeh, L. Mao, P. Stone, J. Biswas, , and X. Xiao. Rethinking social robot navigation: Leveraging the best of two worlds. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2024.
- [31] A. Robicquet, A. Sadeghian, A. Alahi, and S. Savarese. Learning social etiquette: Human trajectory understanding in crowded scenes. In *Proceedings of the European Conference on Computer Vision (ECCV)*, volume 9912, pages 549–565, 10 2016.
- [32] 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.
- [33] A. Rudenko, L. Palmieri, M. Herman, K. M. Kitani, D. M. Gavrila, and K. O. Arras. Human motion trajectory prediction: a survey. *The International Journal of Robotics Research*, 39(8):895–935, 2020.
- [34] S. Saadatnejad, Y. Gao, H. Rezatofighi, and A. Alahi. JRDB-Traj: A dataset and benchmark for trajectory forecasting in crowds, 2023.
- [35] A. Sadeghian, A. Alahi, and S. Savarese. Tracking the untrackable: Learning to track multiple cues with long-term dependencies. *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 300–311, 2017.
- [36] A. Sadeghian, V. Kosaraju, A. Sadeghian, N. Hirose, H. Rezatofighi, and S. Savarese. Sophie: An attentive gan for predicting paths compliant to social and physical constraints. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [37] 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.
- [38] 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.
- [39] S. Samavi, J. R. Han, F. Shkurti, and A. P. Schoellig. SICNav: Safe and interactive crowd navigation using model predictive control and bilevel optimization. *IEEE Transactions on Robotics*, 41:801–818, Jan. 2025.
- [40] 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*, 44(4):568–591, 2025.
- [41] C. Schöller, V. Aravantinos, F. Lay, and A. Knoll. What the constant velocity model can teach us about pedestrian motion prediction. *IEEE Robotics and Automation Letters*, 5(2):1696–1703, 2020.
- [42] 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.
- [43] P. T. Singamaneni, P. Bachiller-Burgos, L. J. Manso, A. Garrell, A. Sanfeliu, A. Spalanzani, and R. Alami. A survey on socially aware robot navigation: Taxonomy and future challenges. *The International Journal of Robotics Research*, 0(0):02783649241230562, 2024.
- [44] A. Stratton, K. Hauser, and C. Mavrogiannis. Characterizing the complexity of social robot navigation scenarios. *IEEE Robotics and Automation Letters*, 10(1):184–191, 2025.
- [45] A. Stratton, P. T. Singamaneni, P. Goyal, R. Alami, and C. Mavrogiannis. How human motion prediction quality shapes social robot navigation performance in constrained spaces. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2026.
- [46] 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*, 44(7):1156–1185, 2025.
- [47] P. Teja S. and R. Alami. HATEB-2: Reactive planning and decision making in human-robot co-navigation. In *Proceedings of the IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, pages 179–186, 2020.
- [48] P. Teja Singamaneni, A. Favier, and R. Alami. Human-aware navigation planner for diverse human-robot interaction contexts. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5817–5824, 2021.
- [49] 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. *International Journal of Robotics Research*, 34(3):335–356, 2015.
- [50] J. van den Berg, S. J. Guy, M. C. Lin, and D. Manocha. Reciprocal n -body collision avoidance. In *Proceedings of the International Symposium on Robotics Research (ISRR)*, pages 3–19, 2009.
- [51] E. Vendrow, D. T. Le, J. Cai, and H. Rezatofighi. JRDB-Pose: A large-scale dataset for multi-person pose estimation and tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [52] M. Walters. *The Design Space for Robot Appearance and Behaviour for Social Robot Companions*. PhD thesis, University of Hertfordshire, Jan. 2008.
- [53] A. Wang, D. Sato, Y. Corzo, S. Simkin, A. Biswas, and A. Steinfeld. Tbd pedestrian data collection: Towards rich, portable, and large-scale natural pedestrian data. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 637–644, 2024.
- [54] R. Wang, W. Wang, and B.-C. Min. Feedback-efficient active preference learning for socially aware robot navigation. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 11336–11343, 2022.
- [55] 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.
- [56] Z. Yan, T. Duckett, and N. Bellotto. Online learning for human classification in 3d lidar-based tracking. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 864–871, 2017.
- [57] 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 ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, page 830–839, 2024.
- [58] Y. Zheng, Y. Yang, K. Mo, J. Li, T. Yu, Y. Liu, C. K. Liu, and L. J. Guibas. GIMO: Gaze-informed human motion prediction in context. In *Proceedings of the European Conference on Computer Vision (ECCV)*, page 676–694, 2022.