

# Following is All You Need: Robot Crowd Navigation Using People as Planners

Yuwen Liao <sup>1</sup>, Graduate Student Member, IEEE, Xinhang Xu <sup>1</sup>, Ruofei Bai <sup>1</sup>, Yizhuo Yang <sup>1</sup>, Muqing Cao <sup>1</sup>, Shenghai Yuan <sup>1</sup>, Member, IEEE, and Lihua Xie <sup>2</sup>, Fellow, IEEE

**Abstract**—Navigating in crowded environments requires the robot to be equipped with high-level reasoning and planning techniques. Existing works focus on developing complex and heavy-weight planners while ignoring the role of human intelligence. Since humans are highly capable agents who are also widely available in a crowd navigation setting, we propose an alternative scheme where the robot utilises people as planners to benefit from their effective planning decisions and social behaviours. Through a set of rule-based evaluations, we identify suitable human leaders who exhibit the potential to guide the robot towards its goal. Using a simple base planner, the robot follows the selected leader through short-horizon subgoals that are designed to be straightforward to achieve. We demonstrate through both simulated and real-world experiments that our novel framework generates safe and efficient robot plans compared to existing planners, even without predictive or data-driven modules. Our method also brings human-like robot behaviours without explicitly defining traffic rules and social norms.

**Index Terms**—Human-aware motion planning, safety in HRI, social HRI.

## I. INTRODUCTION

CROWD navigation is a challenging problem in robotics research as it involves not only static environmental obstacles, but also dynamic agents such as humans. To handle such complex scenarios, recent research focuses on improving the performance of advanced planners with techniques such as predictive planning and reinforcement learning [1], [2], [3], [4]. These works aim to develop intelligent and sophisticated robotic systems that can make correct planning decisions, especially in

Received 26 March 2025; accepted 3 August 2025. Date of publication 13 August 2025; date of current version 20 August 2025. This article was recommended for publication by Associate Editor M. Lagomarsino and Editor A. Peer upon evaluation of the reviewers' comments. This work was supported in part by the Centre for Advanced Robotics Technology Innovation (CARTIN) funded by the National Research Foundation of Singapore through its Medium-Sized Centre Scheme, and in part by the Agency for Science, Technology and Research (A\*STAR) of Singapore under Grant M21K1a0104. (Yuwen Liao and Xinhang Xu contributed equally to this work.) (Corresponding author: Lihua Xie.)

Yuwen Liao, Xinhang Xu, Ruofei Bai, Yizhuo Yang, Shenghai Yuan, and Lihua Xie are with the School of Electrical and Electronic Engineering, Nanyang Technological University, Singapore 639798 (e-mail: yuwen001@ntu.edu.sg; xu0021ng@ntu.edu.sg; ruofei001@ntu.edu.sg; yizhuo001@ntu.edu.sg; shyuan@ntu.edu.sg; elhxie@ntu.edu.sg).

Muqing Cao is with Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213 USA (e-mail: muqingc@andrew.cmu.edu).

Code will be available at <https://github.com/centiLinda/PeopleAsPlanner.git>.

This article has supplementary downloadable material available at <https://doi.org/10.1109/LRA.2025.3598567>, provided by the authors.

Digital Object Identifier 10.1109/LRA.2025.3598567

2377-3766 © 2025 IEEE. All rights reserved, including rights for text and data mining, and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission. See <https://www.ieee.org/publications/rights/index.html> for more information.

©2026 IEEE

Authorized licensed use limited to: Nanyang Technological University Library. Downloaded on August 22, 2025 at 05:06:38 UTC from IEEE Xplore. Restrictions apply.

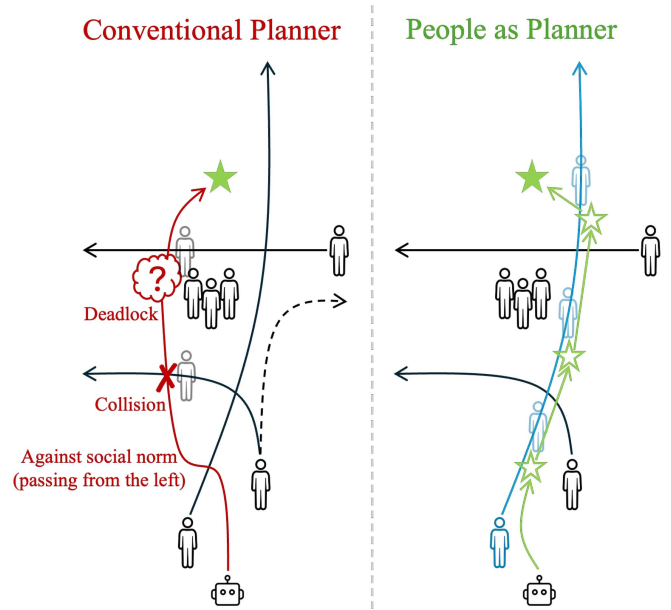


Fig. 1. In a challenging crowd navigation scenario, the robot can easily end up in collisions, deadlocks, and violating implicit social norms. Instead of developing more sophisticated planners, we propose an alternative scheme where the robot follows suitable human leaders to simplify the problem into straightforward subgoal planning.

dense crowds and safety-critical scenarios. On the other hand, humans, perhaps the most intelligent agents in the scene, are solving similar navigation problems as the robot. If we think outside the box, can the robot take advantage of human intelligence to simplify its task?

Inspired by how children follow their parents when walking in crowded streets, we find that following trustworthy leaders is an alternative to planning independently, as illustrated in Fig. 1. We can decompose the crowd navigation problem into high-level decisions and low-level short-horizon movements. The former can be solved using the aforementioned intelligent systems or humans, while the latter only requires simple local planners. Therefore, we propose to rethink human-robot relationships in crowd navigation by handing over high-level decisions to humans. We use people as intelligent planners that can help the robot solve the more difficult part of the crowd navigation problem.

Our proposed method consists of two steps, leader selection and subgoal selection. We start by evaluating nearby humans

based on their observed trajectories, assuming humans are rational and move consistently within a small time window. Through a set of rule-based criteria, we aim to identify suitable human leaders who are 1) ahead of and reachable from the robot's current position, 2) heading towards a similar goal at a desired speed. We then find a subgoal such that it allows the robot to follow the selected leader closely with minimal interruption from other humans. The subgoals are designed to be easily achievable using a simple base planner and we choose Social Force (SF) [5] in our experiments. The leader selection process is performed continuously so the robot can switch between different humans for maximised efficiency. We therefore simplify the task from complex crowd navigation to following suitable human leaders using lightweight local planners, which also enables the robot to demonstrate human-like planning behaviours.

Of relevance to our work is the literature on human-following robots. Many of these works study how to maintain a proper distance from humans in different scenarios to ensure human comfort [6], [7]. Others study how to identify a specific user in complex environments [8], [9], [10]. In these works, a specific user is predefined for the robot to follow without additional context or objectives. Our work is fundamentally different as 1) the robot *actively* chooses a suitable leader to follow, 2) by following humans, we effectively address the challenge of crowd navigation.

Our contributions are threefold. First, we propose a novel crowd navigation framework that utilises human intelligence to decompose the original problem into straightforward subgoal planning. Second, we design a set of evaluation processes to find subgoals near suitable human leaders, which can be effectively reached using a simple base planner. Third, we demonstrate through experiments that the proposed People-as-Planner scheme contributes to safe and efficient robot behaviours in crowded and safety-critical scenarios.

## II. RELATED WORK

### A. Social Robot Navigation

Robot navigation in crowds, or social navigation, has been studied for decades [1], [11]. Traditional model-based methods [5], [12], [13] have been widely adopted in simulations. However, they are less suitable for real-world complex scenarios as they are short-sighted and tend to over-simplify human behaviours [14]. Predictive planning enables the robot to look into the future and avoid potential collisions. They rely heavily on the performance of prediction modules, as inaccurate predictions may lead to accidents and uncertain predictions result in over-conservative behaviours [15], [16]. End-to-end learning-based methods train robot policies using large-scale demonstrations [17], [18], [19]. The performance of the trained networks relies on the accuracy of the simulation platform, or the quality of the collected data [20]. Velocity-field-based methods use crowd motion patterns to inform implicit social norms. Their experiment scenarios are usually simple crowd flows in simulations or indoor settings [21], [22], [23], [24], as complex human movements in the real world are difficult to describe using such homogenised fields. Existing methods

require the robot to plan its path independently. We propose a different approach where the more complex navigation decisions are solved by humans, and the robot simply needs to reach nearby subgoals.

### B. Human-Following Robot

Examples of social navigation scenarios include nursing and warehouse assistance, where robots are tasked to follow a target user. Such systems generally consist of modules including human detection, human tracking, and path planning [25], which have been extensively studied. Recent works focus on more specific tasks such as frontal following [6], sensor fusion [8], [9], [26], and new application scenarios [27]. In terms of following strategy, some works replicate human trajectories exactly [8], [9], while others define subgoals near the leader [25], [26]. We follow the latter as it reduces the perceived intentionality of tracking and results in more natural behaviours. We direct the reader to [25] for a comprehensive review. Existing human-following methods assign a hypothetical task to the robot to follow a predefined human target. In our work, the robot's objective is instead to navigate to its destination by actively selecting and following human leaders.

### C. Human-Robot Relationships

As the robot navigates in crowds, there exist complex relationships between humans and the robot. Apart from considering humans as dynamic obstacles, [28] first proposes the concept of using people as sensors to benefit from human intelligence. This line of work infers the positions of occluded obstacles from observed human behaviours [29], [30], [31], [32]. Inspired by these works, we propose to use people as planners in crowd navigation. There are a few works that also follow human behaviours in robot planning but with different focuses. Refs. [33] and [34] consider human-following as one of the policies when generating trajectory samples in a traditional sampling-based planner. Refs. [8] and [35] utilise human guidance for pathfinding in mapless environments with low human density. Ref. [36] imitates pedestrian behaviours on a narrow sidewalk with limited forms of interactions between humans and the robot. In comparison, our proposed People-as-Planner framework can be applied to various crowded navigation scenarios and is more suitable for real-world robot deployment.

## III. METHODOLOGY

*Problem Definition:* The state of an agent consists of its position and velocity  $\mathbf{x} = [\mathbf{p}^\top, \mathbf{v}^\top]^\top \in \mathbb{R}^4$ . We use subscripts  $H, R$  to represent human and robot, and superscript  $t$  to represent timesteps. We assign a unique ID  $H_i$  where  $i \in \{1, \dots, n\}$  to  $n$  humans that are within the robot's observable range. The robot's goal position is defined as  $\mathbf{p}_g \in \mathbb{R}^2$ .  $\mathbf{X}_H^{0:t}$  represents the historical trajectories of all humans.  $\mathcal{O}$  represents the environmental obstacles within the robot's observable range. In a crowd navigation setting, the robot plan is generated by a planner  $\mathcal{P}$ , denoted as

$$\mathbf{x}_R^{t+1} \leftarrow \mathcal{P}(\mathbf{x}_R^t, \mathbf{x}_H^t, \mathcal{O}, \mathbf{p}_g). \quad (1)$$

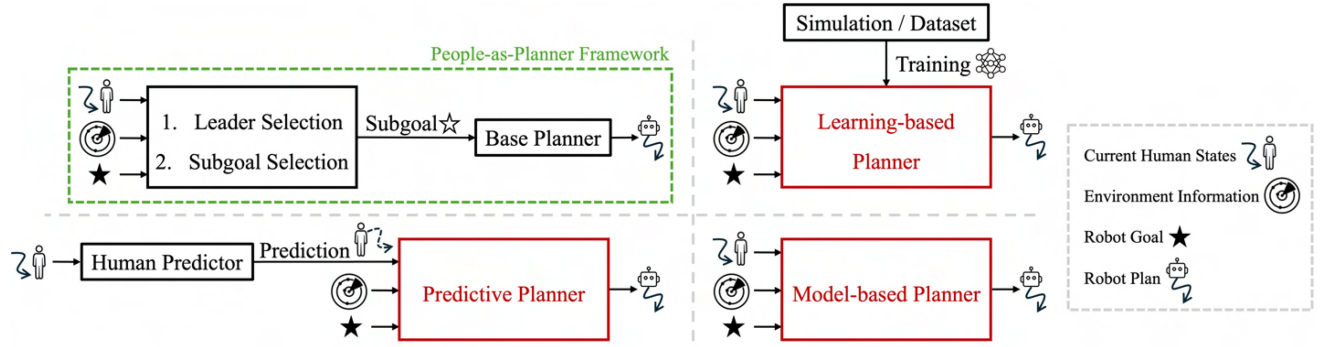


Fig. 2. The proposed people-as-planner framework, compared with common planning methods in crowd navigation. Utilising humans as high-level planners, we only require the same basic inputs as traditional model-based planners, compared to predictive and learning-based planners which require additional modules. We choose SF as the base planner in experiments.

The planner  $\mathcal{P}$  replans at every step until goal is reached.

*Overview:* In our proposed scheme, we first identify a human leader  $H_L^t$  and subsequently define a subgoal  $p_g^t$ :

$$H_L^t = f_{\text{leader}}(\mathbf{x}_R^t, \mathbf{X}_H^{0:t}, \mathcal{O}, \mathbf{p}_g), \quad (2a)$$

$$\mathbf{p}_g^t = f_{\text{subgoal}}(\mathbf{x}_{H_L}^t, \mathbf{x}_H^t), \quad (2b)$$

where  $f_{\text{leader}}(\cdot)$  and  $f_{\text{subgoal}}(\cdot)$  are the leader and subgoal selection processes that will be detailed in Sections III-A and III-B. Robot plan is then generated by base planner  $\mathcal{P}$ :

$$\mathbf{x}_R^{t+1} \leftarrow \mathcal{P}(\mathbf{x}_R^t, \mathbf{x}_H^t, \mathcal{O}, \mathbf{p}_g^t), \quad (3)$$

and the process repeats until goal is reached. We assume that the robot dynamics is incorporated into the planner and our method is planner-agnostic. Since each subgoal  $\mathbf{p}_g^t$  is much easier to reach than the original goal  $\mathbf{p}_g$ , we can use a simple local planner as the base planner. The overall framework is illustrated in Fig. 2.

We will now introduce the two major components of our proposed method: 1) Leader Selection which involves group identification, reachability check, and individual evaluation, 2) Subgoal Selection which defines a straightforward subgoal near the selected leader and adjusts the following speed. Time superscripts are omitted when no confusion is aroused.

### A. Leader Selection

Before evaluating individual humans, we first identify humans that walk closely together as groups. Two Human  $i$  and Human  $j$  are assigned the same group if

$$\|\overrightarrow{\mathbf{p}_j \mathbf{p}_i}\| \leq \tau_{\text{group\_dis}}, \quad \|\overrightarrow{\mathbf{v}_j \mathbf{v}_i}\| \leq \tau_{\text{group\_vel}},$$

where  $\tau_{\text{group\_dis}}$  and  $\tau_{\text{group\_vel}}$  are thresholds for distance and velocity difference. We denote the obtained group information as  $\mathcal{G}$  which includes a list of all groups. Treating groups as a whole can prevent the robot from cutting in between humans that walk together, which respects the social norm.

Next, we evaluate the reachability of each human. As we aim to use simple base planners, the selected leader needs to be directly accessible without the need for complex obstacle avoidance. Spatio-temporal planning [37], [38] is a feasible




Fig. 3. An example of reachability check from robot (red cross) to three candidate leaders (green dot). The purple curves enclose the visible region. Human 4 has a larger score than human 0 for being closer to the robot. Human 3 has a negative score as it is blocked by human 2.

method, but at a high computational cost with prediction uncertainties. Instead, we adopt Line-of-Sight (LoS) distance metric proposed by [39] to define the reachability. As illustrated in Fig. 3, LoS measures the distance from each human to the boundary of the robot's visible region constructed from LiDAR scans. We represent the reachability score as  $f_{\text{reach}}$  and only consider Human  $i$  as a potential leader if

$$f_{\text{reach}}(H_i, \mathbf{x}_R^t, \mathbf{x}_H^t, \mathcal{O}, \mathcal{G}) \geq \tau_{\text{reach}}, \quad (4)$$

where  $\tau_{\text{reach}}$  is the reachability threshold.

We then evaluate each qualified human  $i$  using 3 criteria: 1) average heading, 2) average speed, 3) relative position between the human, the robot, and the robot's goal. Based on  $T$  previous steps, we first evaluate if the human's average heading is towards the robot's goal  $\mathbf{p}_g$ :

$$S_{\text{head}} = \begin{cases} \frac{\tilde{\mathbf{v}}_i \cdot \overrightarrow{\mathbf{p}_i \mathbf{p}_g}}{\|\tilde{\mathbf{v}}_i\| \|\overrightarrow{\mathbf{p}_i \mathbf{p}_g}\|} & \text{if } \frac{\tilde{\mathbf{v}}_i \cdot \overrightarrow{\mathbf{p}_i \mathbf{p}_g}}{\|\tilde{\mathbf{v}}_i\| \|\overrightarrow{\mathbf{p}_i \mathbf{p}_g}\|} \geq \cos \frac{\pi}{4}, \\ -1, & \text{otherwise.} \end{cases} \quad (5)$$

where  $\tilde{\mathbf{v}}_i = \frac{1}{T} \sum_{k=t-T+1}^t \mathbf{v}_i^k$ . We choose not to predict future headings because existing prediction methods may overlook corner cases [40] and increase computational costs.

Second, we compare the human's average speed to the robot's ideal speed  $v_{\text{pref}}$ :

$$S_{\text{vel}} = \begin{cases} \frac{\tilde{v}_i - v_{\text{pref}}}{v_{\text{pref}}}, & \text{if } \tilde{v}_i < v_{\text{pref}}, \\ \max\left(0, 1 - \frac{\tilde{v}_i - v_{\text{pref}}}{v_{\text{pref}}}\right), & \text{otherwise.} \end{cases} \quad (6)$$

where  $\tilde{v}_i = \frac{1}{T} \sum_{k=t-T+1}^t \|\mathbf{v}_i^k\|$ . Humans walking at the ideal speed are given the highest score and the score decreases as human speed exceeds the ideal speed. We penalize slower speeds to encourage active progression.

Third, we compare relative positions to identify humans between robot and its goal, favouring the closer ones:

$$S_{\text{pos}} = \begin{cases} \max\left(0, 1 - \frac{\|\overrightarrow{\mathbf{p}_R \mathbf{p}_i}\|}{r}\right), & \text{if } \frac{\overrightarrow{\mathbf{p}_R \mathbf{p}_i} \cdot \overrightarrow{\mathbf{p}_R \mathbf{p}_g}}{\|\overrightarrow{\mathbf{p}_R \mathbf{p}_i}\| \|\overrightarrow{\mathbf{p}_R \mathbf{p}_g}\|} > 0, \\ -1, & \text{otherwise,} \end{cases} \quad (7)$$

where  $r$  is the robot's observable range.

We can now calculate a weighted score  $S_i$  by

$$S_i = w_{\text{head}} S_{\text{head}} + w_{\text{vel}} S_{\text{vel}} + w_{\text{pos}} S_{\text{pos}}, \quad (8)$$

where  $w_{\text{head}}$ ,  $w_{\text{vel}}$ , and  $w_{\text{pos}}$  are the weights for each score component. If  $S_i$  exceeds a threshold  $\tau_{\text{leader}}$ , Human  $i$  becomes a leader candidate who can lead the robot towards its goal at a desired speed. We add an adjustment term to the previous leader  $S_{H_L^{t-1}}$  to avoid fluctuation between candidates with similar scores. The candidate with the highest score is selected as  $H_L^t$ .

### B. Subgoal Selection

We aim to define a subgoal  $\mathbf{p}_g^t$  that allows the robot to follow  $H_L^t$  through a straightforward path. If  $H_L^t$  belongs to any group in the group list  $\mathcal{G}$ , the group member closest to the robot is used to replace  $H_L^t$  so that the robot will not attempt to interrupt the group.

We first sample a set of position candidates  $\mathbf{p}_m$  between the robot and  $H_L^t$ , defined as

$$\mathbf{p}_m = \mathbf{p}_{H_L^t} - R(\theta_m) \cdot \frac{\overrightarrow{\mathbf{p}_R \mathbf{p}_{H_L^t}}}{\|\overrightarrow{\mathbf{p}_R \mathbf{p}_{H_L^t}}\|} \cdot \frac{d}{\|\overrightarrow{\mathbf{p}_R \mathbf{p}_{H_L^t}}\|}, \quad (9)$$

where  $\theta_m \in \{-\frac{\pi}{4} + m\Delta\theta \mid m = 0, 1, \dots, \lceil \frac{\pi}{2\Delta\theta} \rceil\}$ ;  $d$  is the safe distance from humans; and

$$R(\theta_m) = \begin{bmatrix} \cos(\theta_m) & -\sin(\theta_m) \\ \sin(\theta_m) & \cos(\theta_m) \end{bmatrix}. \quad (10)$$

We select the position furthest from neighbouring humans:

$$\mathbf{p}_g^t = \operatorname{argmax}_{\mathbf{p}_m} \left( \min_{H_i \neq H_L^t} \|\overrightarrow{\mathbf{p}_m \mathbf{p}_{H_i}}\| \right), \quad (11)$$

so to minimize the possibility of collision avoidance.

We adjust the robot's speed in the planner  $\mathcal{P}$  to encourage the robot to follow the leader closely. In our experiments, we update the speed limit  $v_{\text{max}}$  in SF planner directly:

$$v_{\text{max}} = \begin{cases} \|\mathbf{v}_{H_L^t}\|, & \text{if } \|\overrightarrow{\mathbf{p}_R \mathbf{p}_{H_L^t}}\| \leq \tau_{\text{catchup}}, \\ v_{\text{catchup}}, & \text{otherwise,} \end{cases} \quad (12)$$

where  $v_{\text{catchup}}$  is a slightly faster speed when the robot needs to catch up with the leader. For other planners with different settings, this adjustment can be adapted accordingly.

If no  $H_L$  is identified in the previous leader selection step, we will set  $\mathbf{p}_g^t = \mathbf{p}_g$  where it will plan towards the robot's goal directly. We will show in Section IV-C that this happens when there are only a few surrounding humans and using a simple SF planner is sufficient.



Fig. 4. Three test scenes from the SDD dataset: Promenade, crossing, and roundabout. The initial and goal positions are shown in red stars. The direction of travel is interchangeable.

## IV. EXPERIMENTS

### A. Simulation Settings

*Scenarios:* We first perform simulation experiments with real-world trajectory data from the Stanford Drone Dataset (SDD) [41]. SDD was collected mainly from three scenes and we name them Promenade, Crossing, and Roundabout. As shown in Fig. 4, the three scenes exhibit increasing levels of difficulty: 1) **Promenade** has one main road, 2) **Crossing** has a junction with complex traffic, 3) **Roundabout** involves challenging street crossing. For each scene, we select 10 crowded segments from the longest video. The scene densities are listed in Table II. We evaluate our method on all 30 segments and use the densest segment from each scene for comparison study. As we directly playback the dataset, the humans are non-reactive. Reactive human behaviours are studied in the real-world experiments in Section IV-D.

*Implementation Details:* We use a holonomic robot model and the control dynamics is incorporated into the base SF planner. We set  $v_{\text{pref}}$ ,  $r$ ,  $d$  to 1.4m/s [42], 10m, and 0.8m respectively. We set the robot speed limit to 2.0m/s. The rest of the tunable parameters can be found in our code repository. We find the set of parameters through manual tuning and we use the same values for all experiments. Since not all baselines distinguish heterogeneous road users, we set the radius of all agents, including robot, humans, bicycles, and cars to 0.5 m for fair comparisons. For collision evaluation, we provide two sets of agent dimensions: 1) all agents have the same radius of 0.5 m, 2) bicycles and vehicles are given realistic dimensions of 1.9m  $\times$  1m and 4.5m  $\times$  1.9m. We report results for both settings in Section IV-B. Simulations are conducted at 30 Hz on Gazebo platform [43] with AMD 5950XT CPU and RTX3090 graphic card.

*Baselines:* We compare our method with five crowd navigation methods under different frameworks:

- 1) Social Force (SF) [5] calculates attraction from goal and repulsion from obstacles and neighbours.
- 2) Dynamic Window Approach (DWA) [12] searches for the optimal trajectory within a dynamic window that moves towards the goal while avoiding obstacles.
- 3) Optimal Reciprocal Collision Avoidance (ORCA) [13] solves for the optimal action for each agent assuming equal responsibility for avoidance.

TABLE I  
QUANTITATIVE COMPARISONS ON THE DENSEST SEGMENT FROM EACH TEST SCENE

Method	Promenade			Crossing			Roundabout		
	TCC↓	T <sub>avg</sub> (s)↓	D <sub>avg</sub> (m)↓	TCC↓	T <sub>avg</sub> (s)↓	D <sub>avg</sub> (m)↓	TCC↓	T <sub>avg</sub> (s)↓	D <sub>avg</sub> (m)↓
SF [5]	2.66 / 50.13	50.11	62.99	125.46 / 234.13	39.19	54.57	13.42 / 77.21	22.53	31.37
DWA [12]	44.14 / 114.79	48.56	67.03	8.00 / 182.59	29.98	42.43	29.85 / 141.55	27.69	39.16
ORCA [13]	8.00 / 43.11	56.62	74.81	<b>0.00</b> / 80.07	29.70	<b>41.87</b>	2.94 / 53.07	28.41	35.03
Pred2Nav+CV [16]	122.71 / 125.68	48.45	57.34	18.90 / 40.68	32.64	43.50	16.27 / 58.12	19.84	29.37
Pred2Nav+SGAN [16]	138.36 / 139.65	<u>43.92</u>	<u>56.06</u>	71.47 / 117.06	31.98	44.22	23.87 / 83.66	<b>19.40</b>	30.05
HEIGHT [44]	605.68 / 605.68	56.08	<b>43.35</b>	171.45 / 270.46	29.91	44.45	12.17 / 71.76	19.41	<b>26.17</b>
Ours	<b>0.25 / 1.94</b>	<b>40.39</b>	60.48	<b>0.00 / 0.00</b>	<b>29.40</b>	45.10	<b>0.54 / 5.44</b>	19.87	<u>29.17</u>

100 repeated experiments are performed for each segment. For TCC, the first and second values (e.g., 0.25 and 1.94 of our method in the Promenade scene) are calculated by setting all agents with the same radius of 0.5m, and setting bicycles and vehicles with realistic dimensions, respectively. TCC is the frames with collisions under 30Hz simulation (e.g., our method leads to collisions in 0.25 frames in the Promenade scene, averaged over 100 runs). The best performance is in **bold** and the second-best is underlined.

TABLE II  
QUANTITATIVE RESULTS OF OUR METHOD ON THE THREE TEST SCENES WITH 10 SEGMENTS (SEG.) EACH

Seg.	Promenade							Crossing							Roundabout						
	TCC↓	T <sub>avg</sub> (s)↓	D <sub>avg</sub> (m)↓	Ped	Others	Total	Density	TCC↓	T <sub>avg</sub> (s)↓	D <sub>avg</sub> (m)↓	Ped	Others	Total	Density	TCC↓	T <sub>avg</sub> (s)↓	D <sub>avg</sub> (m)↓	Ped	Others	Total	Density
1	1.94	40.39	60.48	21	3	24	high	0.00	29.40	45.10	21	23	44	high	5.44	19.87	29.17	38	51	89	high
2	0.00	40.66	60.27	24	4	28	medium	0.00	28.43	44.73	18	17	35	high	5.33	18.29	29.97	34	49	83	medium
3	0.00	44.40	59.25	13	2	15	low	0.00	28.99	45.48	33	27	60	high	3.62	20.87	28.66	26	43	69	high
4	0.00	37.68	57.16	23	1	24	medium	0.56	30.27	50.50	27	25	52	medium	0.00	16.27	28.45	25	52	77	medium
5	0.00	44.99	60.33	24	2	26	high	12.41	34.04	48.09	27	24	51	medium	0.00	17.70	26.60	15	48	63	low
6	1.73	44.33	61.60	24	3	27	high	0.00	27.26	44.17	21	17	38	medium	0.00	16.70	25.99	6	29	35	low
7	3.10	42.72	59.66	23	4	27	high	6.70	32.45	46.93	12	17	29	medium	8.38	25.08	34.54	33	46	79	high
8	0.00	35.95	58.66	16	3	19	low	5.00	27.82	46.57	19	18	37	medium	1.75	18.53	26.81	15	47	62	low
9	14.62	42.87	56.94	18	4	22	medium	12.69	31.13	47.86	26	22	48	medium	7.64	20.11	30.13	24	51	75	medium
10	0.00	36.88	56.09	21	2	23	medium	0.00	35.50	46.92	27	19	46	low	7.56	17.45	27.77	28	51	79	medium
Avg	2.14	41.09	59.04	\\	\\	\\	\\	3.74	30.53	46.64	\\	\\	\\	\\	3.97	19.09	28.80	\\	\\	\\	\\
STD	4.53	3.32	1.80	\\	\\	\\	\\	5.22	2.73	1.89	\\	\\	\\	\\	3.37	2.58	2.46	\\	\\	\\	\\

100 repeated experiments are performed for each segment. For TCC, we use realistic dimensions for bicycles and vehicles, which corresponds to the second value of TCC in Table I. We count the total number of road users in each segment, including pedestrians, bicycles, and vehicles. We record the crowd density along the robot's path, determined through visual inspection.

- Pred2Nav** [16] incorporates human trajectory predictors into a sampling-based MPC. We compare two predictors implemented by the author, constant velocity (CV) and SGAN [45] which is a neural network.
- HEIGHT** [44] is an end-to-end reinforcement learning method that considers crowd interactions. We follow the training procedures from the provided repository to train the model for holonomic robots.

*Evaluation Metrics:* We perform 100 repeated runs for each segment and evaluate the average performance to count for the uncertainty of physics engine approximation in simulation. We mark it as a collision if the robot enters within the agent dimensions. The following three metrics are used:

- Total Collision Count (TCC) is the number of frames where collision occurs. Experiment continues after a collision so there could be multiple counts in one run.
- Average Time (T<sub>avg</sub>) taken to reach the goal.
- Average Distance (D<sub>avg</sub>) taken to reach the goal.

## B. Quantitative Evaluation

As shown in Table I, our method shows the best safety awareness with high efficiency across all three scenes. Our method allows the robot to adjust its speed according to the leader, so it can reach the goal fast even when additional manoeuvres are required for effective collision avoidance. Notice that although we set the radius for all agents to 0.5 m during the experiments, when we calculate collisions using real dimensions for other road users (namely bicycles and vehicles), our method maintains a low TCC, while the values increase dramatically for all other

baselines. This implies that our method can keep a larger distance from dangerous road users without explicitly considering agent types in the inputs, which benefits from the human leader's planning decisions.

In comparison, ORCA tends to get too close to humans since it assumes reciprocal avoidance behaviours, which lead to more collisions than ours. SF and DWA are less robust in highly dynamic environments, as shown by the inconsistent performance across the three scenes. For Pred2Nav, although SGAN gives more accurate human predictions on large-scale datasets [45], it doesn't consistently improve the planning performance. HEIGHT is less sensitive to collisions because it is not explicitly enforced like the other methods. It demonstrates biased behaviours such as moving in a curved path, which results in heavy collisions with the wall in the Promenade scene. We also show that by simplifying complex crowd navigation problems using people as planners, we can achieve the best performance using a simple base planner SF, which is less capable when planning independently.

We summarise the types of information and procedures required to deploy each method in Table III. All methods except DWA require current states of humans to consider human avoidance and interaction explicitly. All baselines except Pred2Nav require environmental information such as obstacle maps and point clouds for obstacle avoidance. Although Pred2Nav does not consider static obstacles, none of its TCC reported in Table I comes from collision with walls. Our method uses point clouds for human reachability evaluations. In addition, Pred2Nav requires a trajectory prediction module to obtain future human states. HEIGHT, as well as SGAN used by Pred2Nav, require

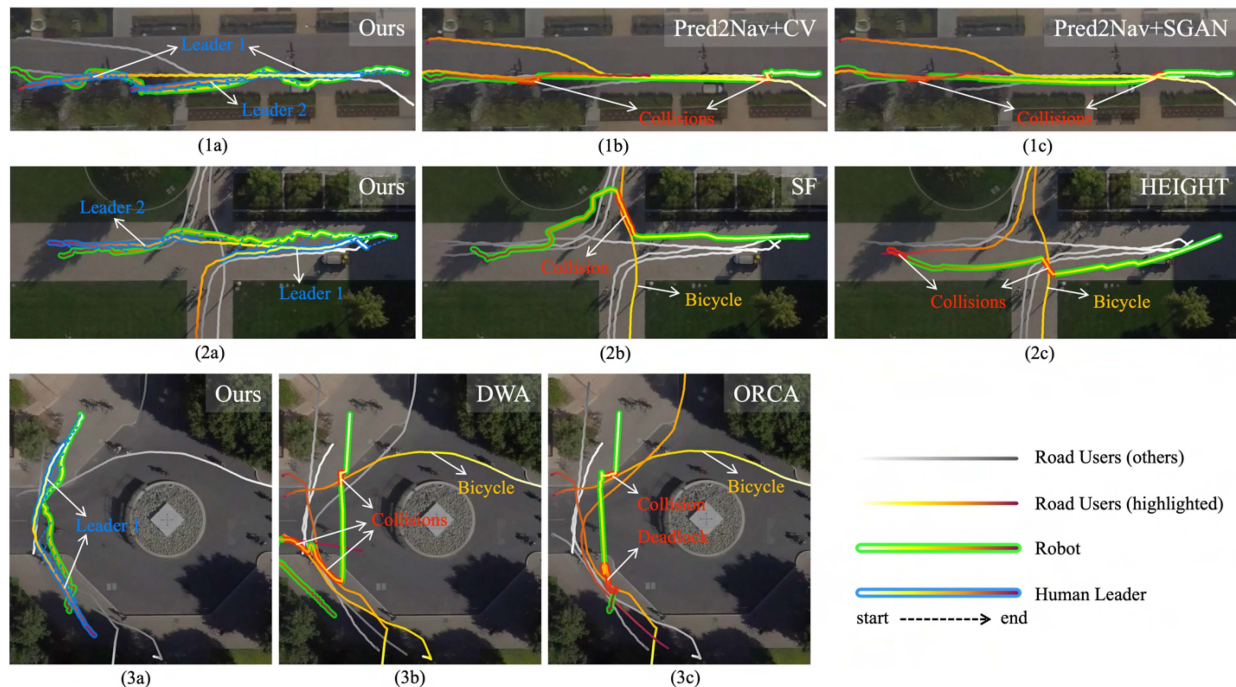


Fig. 5. Qualitative comparisons in simulation experiments. The three rows correspond to the promenade, crossing, and roundabout scenes respectively. The temporal progression is represented by the gradient colour (from light to dark) of each trajectory. The robot and the selected leaders are highlighted by different glow colours. For simplicity, we select a few agents to draw for each scene to highlight the interactions. The full experiment recordings can be found in the supplementary video.

TABLE III  
COMPARISON OF DEPLOYMENT REQUIREMENTS BETWEEN BASELINES AND OUR METHOD

Method	Current human states	Future human states	Environment information	Training
SF [5]	✓	×	✓	×
DWA [12]	×	×	✓	×
ORCA [13]	✓	×	✓	×
Pred2Nav+CV [16]	✓	✓	×	×
Pred2Nav+SGAN [16]	✓	✓	×	✓
HEIGHT [44]	✓	×	✓	✓
Ours	✓	×	✓	×

training on either simulations or large-scale datasets. We can see that our method only needs the basic inputs to achieve the best overall performance. Our method is further evaluated over all 30 segments from the three scenes, as shown in Table II. We demonstrate strong robustness as the performance is consistent under different traffic conditions and crowd densities.

### C. Qualitative Evaluation

Using people as planners, our method demonstrates efficient and socially-aware behaviours that contribute to the outstanding results presented above. For the Promenade scene in Fig. 5(1a), the robot switched to the second leader when the first leader was temporarily occluded, and switched back when the second leader stopped moving. In Fig. 5(1b), Pred2Nav resulted in collisions twice. From the sharp turns before the collisions, we can infer that the robot tried to avoid the humans, but it acted too late.

Compared to Fig. 5(1c), predicting with SGAN did not bring much benefit. For the Crossing scene, both SF and HEIGHT resulted in collisions with a fast-moving bicycle, as shown in Fig. 5(2b) and (2c). Our method, as shown in Fig. 5(2a), avoided this encounter by switching between the two leaders.

For the Roundabout scene, DWA first collided with a bicycle, then with two humans walking face-to-face, as shown in Fig. 5(3b). In Fig. 5(3c), ORCA only collided with the bicycle, but was forced to stop by the two humans and only continued moving after the humans walked by. In comparison, as shown in Fig. 5(3a), our method followed the leader to cross the road safely, switched to default planner in the middle when the leader was temporarily occluded, and then continued to follow towards the goal. An interesting observation in this scene is that, although traffic rules are not explicitly defined, our method follows the human to cross the road from the appropriate point on the sidewalk, which avoids a dangerous encounter with the bicycle. Other baselines, however, move towards the goal directly, which inevitably increases the difficulty of avoiding fast-moving road users on the main road. This makes our method especially useful in real-world deployments, as it is difficult to define every rule in every scenario, our method can simply follow the surrounding agents to respect the “unknown rules”.

### D. Real-World Experiment

In addition to the performance comparison against baselines using recorded human data, we further demonstrate the effectiveness of our method when interacting with real-world reactive




Fig. 6. Events observed in the real-world experiment. Green arrows represent robot subgoal planning. Blue colours represent selected leaders and their movements. Grey arrows represent movements of other humans. The dashed circle in event 1 represents the previous leader.

humans. We conduct the experiment in a 150 m crowded campus corridor with over 40 human encounters, which is similar to the previous SDD scenes in terms of environment and human density. We use a 2D LiDAR to obtain point clouds for reachability check, and use a depth camera for human tracking. To ensure human safety and privacy, we reduce the robot speed limit to 1.0 m/s, stop the robot when getting too close to pedestrians, and do not collect their identifiable information during experiments.

In Fig. 6, we illustrate four events, most of which have been observed in the simulation experiments. In the first event, the robot switched to a new leader when the previous leader slowed down. This ensures the robot continues to progress efficiently without being stuck behind static humans. In the second and fourth events, the robot moved through a complex crowd and a narrow gate without deadlocks. In the third event, the robot followed the social norm to keep to the left, which avoided encounters with humans walking from the opposite direction. The robot benefits from the proposed people-as-planner scheme and makes intelligent decisions in real-world crowded environments. Our method is robust to disturbances such as inaccurate human tracking, hardware delays, etc., and successfully completes the run without collisions. The complete experiment can be found in the supplementary video.

### E. Limitations

During simulation experiments, we found that our method may fail when interacting with aggressive road users, as shown in Fig. 7. Future work could incorporate data-driven policies trained on real-world corner cases. Some existing works allow the robot to actively influence human behaviours [46], [47], [48], which may also be useful, especially in difficult scenarios such as moving against a dense crowd.

The dynamic subgoal setting in the proposed framework may potentially affect trajectory smoothness. In the current experiments, our method does not show significant fluctuations compared to the baselines (comparison results on trajectory smoothness can be found in the supplementary video). Nevertheless, to improve the smoothness, future work may consider other types of robot dynamics apart from the holonomic model, or incorporate various trajectory smoothing techniques [49].

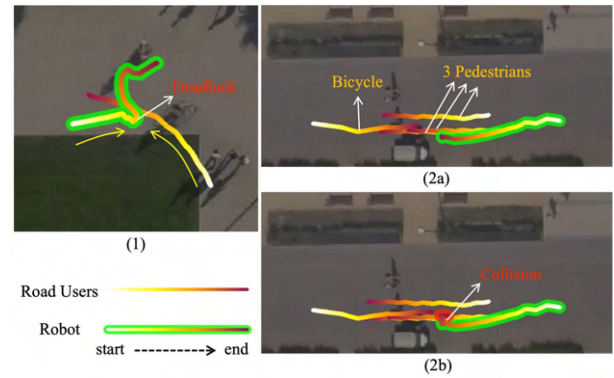


Fig. 7. Failure cases in simulation experiments. In (1), the robot was confused by a human blocking the way and forced into a short deadlock. In (2a), the robot was following a group and did not detect an oncoming bicycle due to occlusion. A dataset labelling error caused the bicycle to unrealistically run over the pedestrians, leading to a sudden collision with the robot in (2b).

## V. CONCLUSION

This paper presents a novel crowd navigation scheme which utilises people as external intelligent planners. The more complex planning decisions are handed over to humans so the robot's task is simplified into straightforward subgoal planning. The experiments demonstrate robot behaviours that are safe, efficient, and socially compliant in various crowded environments. Our proposed framework offers new insights into human-robot relationships in social navigation.

## REFERENCES

- [1] C. Mavrogiannis et al., "Core challenges of social robot navigation: A survey," *ACM Trans. Hum.-Robot Interaction*, vol. 12, no. 3, pp. 1–39, 2023.
- [2] R. Korbmayer and A. Tordeux, "Review of pedestrian trajectory prediction methods: Comparing deep learning and knowledge-based approaches," *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 12, pp. 24126–24144, Dec. 2022.
- [3] 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 *Proc. 2023 IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2023, pp. 7442–7447.
- [4] J. Guo, G. Zhou, H. Huang, and C. Huang, "Advancements in UAV path planning: A deep reinforcement learning approach with soft actor-critic for enhanced navigation," *Unmanned Syst.*, vol. 13, no. 4, pp. 1065–1084, 2025.
- [5] D. Helbing and P. Molnar, "Social force model for pedestrian dynamics," *Phys. Rev. E*, vol. 51, no. 5, 1995, Art. no. 4282.
- [6] S. Leisiazar, S. R. R. Rohani, E. J. Park, A. Lim, and M. Chen, "Adapting to frequent human direction changes in autonomous frontal following robots," *IEEE Robot. Automat. Lett.*, vol. 10, no. 3, pp. 2934–2941, Mar. 2025.
- [7] C.-P. Lam, C.-T. Chou, K.-H. Chiang, and L.-C. Fu, "Human-centered robot navigation—towards a harmoniously human–robot coexisting environment," *IEEE Trans. Robot.*, vol. 27, no. 1, pp. 99–112, Feb. 2011.
- [8] A. Antonucci, P. Bevilacqua, S. Leonardi, L. Paolopoli, and D. Fontanelli, "Humans as path-finders for mobile robots using teach-by-showing navigation," *Auton. Robots*, vol. 47, no. 8, pp. 1255–1273, 2023.
- [9] R. Algabri and M.-T. Choi, "Deep-learning-based indoor human following of mobile robot using color feature," *Sensors*, vol. 20, no. 9, 2020, Art. no. 2699.
- [10] M. Gupta, S. Kumar, L. Behera, and V. K. Subramanian, "A novel vision-based tracking algorithm for a human-following mobile robot," *IEEE Trans. Syst., Man, Cybern. Syst.*, vol. 47, no. 7, pp. 1415–1427, Jul. 2017.

- [11] M. Hoy, A. S. Matveev, and A. V. Savkin, "Algorithms for collision-free navigation of mobile robots in complex cluttered environments: A survey," *Robotica*, vol. 33, no. 3, pp. 463–497, 2015.
- [12] D. Fox, W. Burgard, and S. Thrun, "The dynamic window approach to collision avoidance," *IEEE Robot. Automat. Mag.*, vol. 4, no. 1, pp. 23–33, Mar. 1997.
- [13] J. V. D. Berg, S. J. Guy, M. Lin, and D. Manocha, "Reciprocal n-body collision avoidance," in *Proc. Robot. Res.: 14th Int. Symp.*, 2011, pp. 3–19.
- [14] P. Trautman and A. Krause, "Unfreezing the robot: Navigation in dense, interacting crowds," in *Proc. 2010 IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2010, pp. 797–803.
- [15] K. Ryu and N. Mehr, "Integrating predictive motion uncertainties with distributionally robust risk-aware control for safe robot navigation in crowds," May 2024, pp. 2410–2417, doi: [10.1109/ICRA57147.2024.10610404](https://doi.org/10.1109/ICRA57147.2024.10610404).
- [16] S. Poddar, C. Mavrogiannis, and S. S. Srinivasa, "From crowd motion prediction to robot navigation in crowds," in *Proc. 2023 IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2023, pp. 6765–6772.
- [17] Z. Xie and P. Dames, "DRL-VO: Learning to navigate through crowded dynamic scenes using velocity obstacles," *IEEE Trans. Robot.*, vol. 39, no. 4, pp. 2700–2719, Aug. 2023.
- [18] Y. Yang, J. Jiang, J. Zhang, J. Huang, and M. Gao, "ST<sup>2</sup>: Spatial-temporal state transformer for crowd-aware autonomous navigation," *IEEE Robot. Automat. Lett.*, vol. 8, no. 2, pp. 912–919, Feb. 2023.
- [19] H. Karnan, G. Warnell, X. Xiao, and P. Stone, "VOILA: Visual-observation-only imitation learning for autonomous navigation," in *Proc. 2022 Int. Conf. Robot. Automat.*, 2022, pp. 2497–2503.
- [20] K. Zhu and T. Zhang, "Deep reinforcement learning based mobile robot navigation: A review," *Tsinghua Sci. Technol.*, vol. 26, no. 5, pp. 674–691, Oct. 2021.
- [21] D. Dugas, K. Cai, O. Andersson, N. Lawrance, R. Siegwart, and J. J. Chung, "FlowBot: Flow-based modeling for robot navigation," in *Proc. 2022 IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2022, pp. 8799–8805.
- [22] K. Cai, W. Chen, D. Dugas, R. Siegwart, and J. J. Chung, "Sampling-based path planning in highly dynamic and crowded pedestrian flow," *IEEE Trans. Intell. Transp. Syst.*, vol. 24, no. 12, pp. 14732–14742, Dec. 2023.
- [23] Y. Zhu, A. Rudenko, L. Palmieri, L. Heuer, A. J. Lilienthal, and M. Magnusson, "Fast online learning of CLiFF-maps in changing environments," 2024, *arXiv:2410.12237*.
- [24] Y. Zhu et al., "CLiFF-LHMP: Using spatial dynamics patterns for long-term human motion prediction," in *Proc. 2023 IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2023, pp. 3795–3802.
- [25] M. J. Islam, J. Hong, and J. Sattar, "Person-following by autonomous robots: A categorical overview," *Int. J. Robot. Res.*, vol. 38, no. 14, pp. 1581–1618, 2019.
- [26] C. Scheidemann et al., "Obstacle-avoidant leader following with a quadruped robot," 2024, *arXiv:2410.00572*.
- [27] W.-C. Chen, C.-L. Lin, Y.-Y. Chen, and H.-H. Cheng, "Quadcopter drone for vision-based autonomous target following," *Aerospace*, vol. 10, no. 1, 2023, Art. no. 82.
- [28] M. Lewis, H. Wang, P. Velagapudi, P. Scerri, and K. Sycara, "Using humans as sensors in robotic search," in *Proc. 12th Int. Conf. Inf. Fusion*, 2009, pp. 1249–1256.
- [29] O. Afolabi, K. Driggs-Campbell, R. Dong, M. J. Kochenderfer, and S. S. Sastry, "People as sensors: Imputing maps from human actions," in *Proc. 2018 IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2018, pp. 2342–2348.
- [30] M. Itkina, Y.-J. Mun, K. Driggs-Campbell, and M. J. Kochenderfer, "Multi-agent variational occlusion inference using people as sensors," in *Proc. 2022 Int. Conf. Robot. Automat.*, 2022, pp. 4585–4591.
- [31] Y.-J. Mun, M. Itkina, S. Liu, and K. Driggs-Campbell, "Occlusion-aware crowd navigation using people as sensors," in *Proc. 2023 IEEE Int. Conf. Robot. Automat.*, 2023, pp. 12031–12037.
- [32] T. Qiu and D. Fridovich-Keil, "Inferring occluded agent behavior in dynamic games from noise corrupted observations," *IEEE Robot. Automat. Lett.*, vol. 9, no. 12, pp. 11489–11496, Dec. 2024.
- [33] D. Mehta, G. Ferrer, and E. Olson, "Autonomous navigation in dynamic social environments using multi-policy decision making," in *Proc. 2016 IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2016, pp. 1190–1197.
- [34] Y. Chen and Y. Lou, "A unified multiple-motion-mode framework for socially compliant navigation in dense crowds," *IEEE Trans. Automat. Sci. Eng.*, vol. 19, no. 4, pp. 3536–3548, Oct. 2022.
- [35] S. Buckeridge, P. Careno-Medrano, A. Cosgun, E. Croft, and W. P. Chan, "Mapless urban robot navigation by following pedestrians," in *Proc. 2023 IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2023, pp. 6787–6792.
- [36] Y. Du et al., "Group Surfing: A pedestrian-based approach to side-walk robot navigation," in *Proc. 2019 Int. Conf. Robot. Automat.*, 2019, pp. 6518–6524.
- [37] W. Ding, L. Zhang, J. Chen, and S. Shen, "Safe trajectory generation for complex urban environments using spatio-temporal semantic corridor," *IEEE Robot. Automat. Lett.*, vol. 4, no. 3, pp. 2997–3004, Jul. 2019.
- [38] T. Zhang, M. Fu, W. Song, Y. Yang, and M. Wang, "Trajectory planning based on spatio-temporal map with collision avoidance guaranteed by safety strip," *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 2, pp. 1030–1043, Feb. 2022.
- [39] R. Bai, S. Yuan, K. Li, H. Guo, W.-Y. Yau, and L. Xie, "Realm: Real-time line-of-sight maintenance in multi-robot navigation with unknown obstacles," 2025, *arXiv:2502.15162*.
- [40] N. A. Madjid et al., "Trajectory prediction for autonomous driving: Progress, limitations, and future directions," 2025, *arXiv:2503.03262*.
- [41] A. Robicquet, A. Sadeghian, A. Alahi, and S. Savarese, "Learning social etiquette: Human trajectory understanding in crowded scenes," in *Proc. Comput. Vis.—ECCV 2016: 14th Eur. Conf.*, 2016, pp. 549–565.
- [42] R. V. Levine and A. Norenzayan, "The pace of life in 31 countries," *J. Cross-Cultural Psychol.*, vol. 30, no. 2, pp. 178–205, 1999.
- [43] N. Koenig and A. Howard, "Design and use paradigms for Gazebo, an open-source multi-robot simulator," in *Proc. 2004 IEEE/RSJ Int. Conf. Intell. robots Syst.*, 2004, vol. 3, pp. 2149–2154.
- [44] S. Liu, H. Xia, F. C. Pouria, K. Hong, N. Chakraborty, and K. Driggs-Campbell, "HEIGHT: Heterogeneous interaction graph transformer for robot navigation in crowded and constrained environments," 2024, *arXiv:2411.12150*.
- [45] A. Gupta, J. Johnson, L. Fei-Fei, S. Savarese, and A. Alahi, "Social GAN: Socially acceptable trajectories with generative adversarial networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 2255–2264.
- [46] D. Dugas, J. Nieto, R. Siegwart, and J. J. Chung, "IAN: Multi-behavior navigation planning for robots in real, crowded environments," in *Proc. 2020 IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2020, pp. 11368–11375.
- [47] Y. Che, A. M. Okamura, and D. Sadigh, "Efficient and trustworthy social navigation via explicit and implicit robot–human communication," *IEEE Trans. Robot.*, vol. 36, no. 3, pp. 692–707, Jun. 2020.
- [48] G. Angelopoulos, A. Rossi, C. D. Napoli, and S. Rossi, "You are in my way: Non-verbal social cues for legible robot navigation behaviors," in *Proc. 2022 IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2022, pp. 657–662.
- [49] A. Ravankar, A. A. Ravankar, Y. Kobayashi, Y. Hoshino, and C.-C. Peng, "Path smoothing techniques in robot navigation: State-of-the-art, current and future challenges," *Sensors*, vol. 18, no. 9, 2018, Art. no. 3170.