From Cognition to Precognition: A Future-Aware Framework for Social Navigation

Zeying Gong[1], Tianshuai Hu[2], Ronghe Qiu[1] and Junwei Liang[1,2*]

***Abstract*— To navigate safely and efficiently in crowded spaces, robots should not only perceive the current state of the environment but also anticipate future human movements. In this paper, we propose a reinforcement learning architecture, namely *Falcon*, to tackle socially-aware navigation by explicitly predicting human trajectories and penalizing actions that block future human paths. To facilitate realistic evaluation, we introduce a novel SocialNav benchmark containing two new datasets, Social-HM3D and Social-MP3D. This benchmark offers large-scale photo-realistic indoor scenes populated with a reasonable amount of human agents based on scene area size, incorporating natural human movements and trajectory patterns. We conduct a detailed experimental analysis with the state-of-the-art learning-based method and two classic rule-based path-planning algorithms on the new benchmark. The results demonstrate the importance of future prediction and our method achieves the best task success rate of 55% while maintaining about 90% personal space compliance. We have already released our code and datasets.**

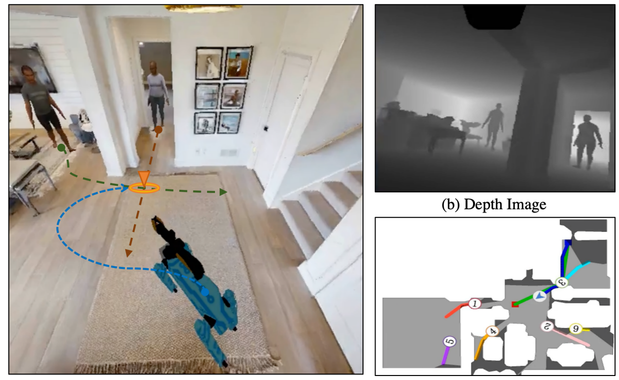

(a) Third-person RGB Perspective (b) Depth Image (c) Top-down Map

Fig. 1: We integrate trajectory prediction into the SocialNav task. In (a), the robot navigates toward a goal while predicting human trajectories (dashed lines) and avoiding them. The robot uses depth input as shown in (b). (c) offers a top-down map for reference.

I. INTRODUCTION

Social navigation (SocialNav) refers to autonomous robots adhering to *social norms* and *social etiquette* while navigating human-shared environments [1]. This task challenges traditional visual navigation methods, particularly in human-populated settings where collision avoidance is critical.

Existing RL-based approaches [2], [3], [4] often struggle due to limited foresight and reliance on global information [5], [6], [7]. For instance, consider Fig. 1, where a robot navigates toward a goal intersecting two humans' future paths. Traditional methods may fail due to limited foresight or over-reliance on global data. In contrast, our approach predicts human trajectories explicitly, enabling social compliance and long-term collision avoidance.

Human trajectory forecasting improves navigation in dynamic environments [8], [9], [10], but it is primarily applied in outdoor scenarios like autonomous driving [11], [12], [13]. Motivated by this, we integrate trajectory prediction into SocialNav, addressing indoor challenges such as limited space and maneuverability [7].

To address these challenges, we propose ***Falcon***, a future-aware SocialNav framework with precognition. *Falcon*'s key novelties include: 1) the **Social Cognition Penalty**, which penalizes trajectory obstructions, promoting proactive collision avoidance; and 2) the **Spatial-Temporal Precognition Module**, incorporating socially-aware auxiliary tasks like trajectory prediction to enhance future dynamics understanding during training.

Another challenge in SocialNav is unrealistic configurations [14]. Current methods oversimplify environments, neglecting scene complexity [15], [16], [17], and assume access to global information [18]. To address this, we introduce a novel SocialNav benchmark with two datasets, **Social-HM3D** and **Social-MP3D**. Built using 3D-reconstructed real-world indoor scenes, these datasets feature realistic human movements and robot animations. The robot uses only egocentric inputs and a point goal during inference, without relying on global maps or known trajectories. Our benchmark provides a more realistic representation of SocialNav. *Falcon* achieves state-of-the-art performance on this benchmark, with a 55% success rate and high social compliance.

Overall, our main contributions are as follows:

- We introduce the first realistic SocialNav benchmark with two novel datasets, **Social-HM3D** and **Social-MP3D**, featuring large-scale photo-realistic scenes with realistic human and robot animations.
- We propose a new and effective SocialNav framework, ***Falcon***, that integrates explicit trajectory prediction, allowing the robot to perceive and predict for safe, socially comfortable and effective navigation.
- We establish a new state-of-the-art result compared to prior approaches on the proposed benchmark.

[1] The Hong Kong University of Science and Technology (Guangzhou). {zgong313,rqiu683}@connect.hkust-gz.edu.cn, junweiliang@hkust-gz.edu.cn * Corresponding author.
[2] The Hong Kong University of Science and Technology. thuaj@connect.ust.hk

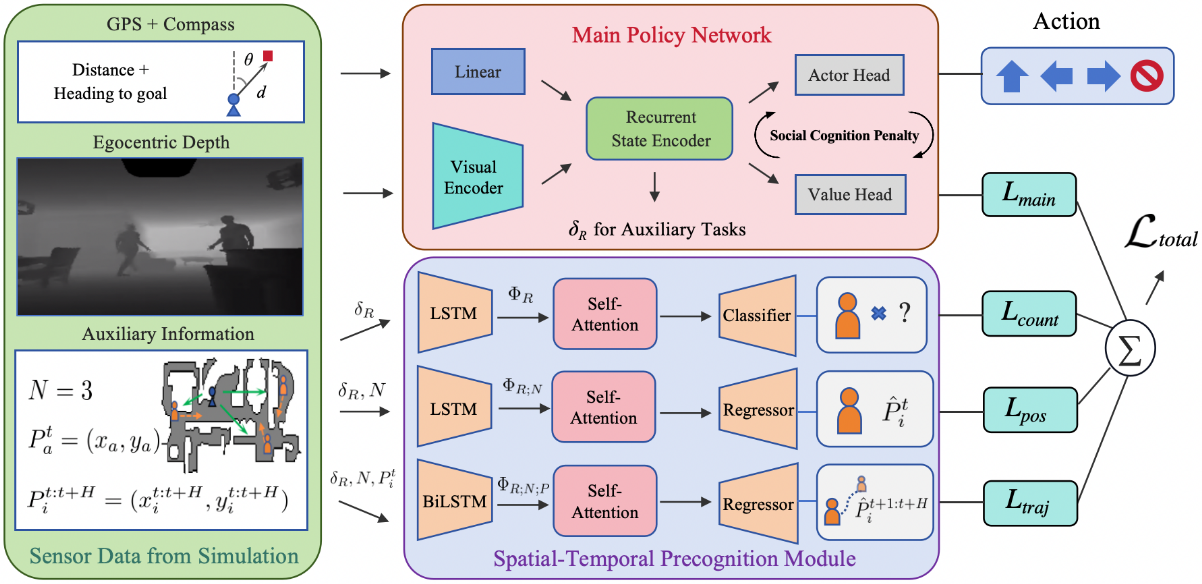

Fig. 2: ***Falcon* Overview**: The main policy network (top-right) takes Depth and GPS+Compass data as input. Its behavior is guided by the **Social Cognition Penalty**, which encourages socially compliant navigation and generates the main loss. During training, the output of the network's state encoder, combined with auxiliary information from the Habitat simulator, is processed by the **Spatial-Temporal Precognition Module** (bottom-right). Three socially-aware auxiliary tasks are then performed, producing auxiliary losses. The total loss is computed by weighting the main loss with the auxiliary losses.

II. Related Works

A. Social Navigation.

In this paper, we focus on the SocialNav task [7], [14], introduced in the iGibson SocialNav Challenge [19]. This task extends PointGoal Navigation (PointNav) by adding moving humans, though the humans in the challenge have unrealistic movements. Our work uses the Habitat 3.0 simulator [20] to introduce realistic human movements and animations. SocialNav has been widely studied in robotics, computer vision, and social behavior analysis [21], [22]. Research in collision-free multi-agent navigation [23], [24], [25] and dynamic environments [26] has addressed challenges posed by human presence [27], [28], [29], [30], [31]. Our method introduces explicit trajectory prediction within auxiliary tasks to train an RL-based agent for SocialNav.

B. Human Trajectory Prediction.

Human trajectory prediction is vital for enabling safe and intelligent behavior in autonomous systems [32], [33]. Traditional approaches rely on physical models like the Social Force model [28], which uses forces to simulate social behaviors. Methods fall into three categories: physics-based models [34], [35], [36], learning-based methods [37], [38], [39], and planning-based methods [40], [41], [42]. Our approach integrates socially-aware information into the agent's navigation policy in dynamic scenes.

III. Methodology

A. Problem Formulation

We consider a social navigation task where a robot navigates in an environment with N humans. Starting from an initial configuration, the robot aims to reach a goal while avoiding collisions with static obstacles and dynamic humans. The objective is:

$$\tau_a = \arg\min_{\tau \in T}(c_a(\tau) + \lambda_a c_{sa}(\tau, \tau_{1:N})) \tag{1}$$

where c_a is the path cost, c_{sa} accounts for social norms, and λ_a is a weight factor.

B. Overview of Falcon

Figure 2 shows the architecture of Falcon. The Main Policy Network takes depth images and point goals as inputs, outputs actions using DD-PPO, and is guided by a Social Cognition Penalty (SCP). During training, the Spatial-Temporal Precognition Module performs auxiliary tasks to enhance the agent's understanding of future dynamics.

C. Main Policy Network

The main policy network consists of a ResNet-50 visual encoder, an LSTM for temporal features, and actor-critic heads for action prediction and reward estimation. The reward function combines PointNav rewards with SCP penalties:

$$R_{socialnav} = R_{pointnav} - R_{scp} \tag{2}$$

where R_{scp} includes penalties for obstacle collisions, human proximity, and trajectory obstructions.

D. Spatial-Temporal Precognition Module

This module performs three socially-aware auxiliary tasks:

Human Count Estimation. Predicts the number of humans using a classifier:

$$L_{count} = -\sum_{k=0}^{M} n_k \log(\hat{n}_k) \tag{3}$$

Dataset	Num. Scenes	Scene Type	Max Num. Humans	Natural Motions
iGibson-SN [43]	15	residence	3	✗
Isaac Sim [44]	7	residence, office, depot, etc.	7	✓
HabiCrowd [45]	480	residence, office, gym, etc.	40	✗
HM3D-S [18]	900	residence, office, shop, etc.	3	✗
Social-HM3D	844	residence, office, shop, etc.	6	✓
Social-MP3D	72	residence, office, gym, etc.	6	✓

TABLE I: Statistics Comparison of SocialNav Datasets/Simulators: Our proposed Social-HM3D and Social-MP3D datasets feature extensive scene diversity and realistic interaction design, addressing the shortcomings of previous datasets which often relied on over-simplified human behaviors and imbalanced interaction dynamics.

where n_k is the true count and $\hat{n}_k$ is the predicted probability.

Current Position Tracking. Tracks human positions relative to the robot using regression:

$$L_{pos} = \frac{1}{|M|} \sum_{i \in M} \|\hat{P}_i^t - P_i^t\|^2 \tag{4}$$

where $\hat{P}_i^t$ is the predicted position and P_i^t is the true position.

Future Trajectory Forecasting. Predicts human trajectories over multiple time steps:

$$L_{traj} = \frac{1}{|M|} \sum_{i \in M} \|\hat{P}_i^{t+1:t+H} - P_i^{t+1:t+H}\|^2 \tag{5}$$

where $\hat{P}_i^{t+1:t+H}$ is the predicted trajectory and $P_i^{t+1:t+H}$ is the ground truth.

The total loss is computed as:

$$L_{total} = \beta_{main} L_{main} + \beta_{aux} L_{aux} \tag{6}$$

where $L_{aux} = L_{count} + L_{pos} + L_{traj}$.

IV. EXPERIMENTS

A. Datasets

Existing SocialNav datasets[43], [44] often lack scene diversity or realistic human behaviors. Others [45], [18], while rich in scenes, fail to balance human density or provide natural movement. To address these limitations, we introduce a benchmark with two simulation datasets: **Social-HM3D** and **Social-MP3D**, derived from HM3D[46] and MP3D [47].

Table I compares our datasets with existing ones. Our datasets feature diverse environments, realistic human motions, and balanced interaction dynamics, enabling more effective social navigation algorithm development.

Realistic Human Behaviors. Our dataset provides goal-driven human trajectories with natural movement patterns, unlike random walks or repetitive movements [43], [45]. Each human alternates between moving towards two goals and resting, stopping once tasks are complete. Their walking speeds vary between 0.8–1.2 times the robot's speed. Collisions are avoided using ORCA [23], and realistic animations enhance visual authenticity.

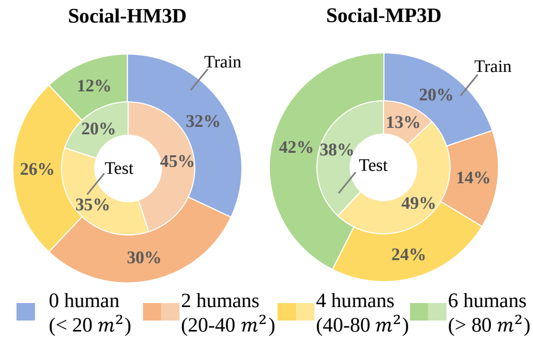

Fig. 3: Human Distribution by Scene Area in Social-HM3D and Social-MP3D (Train/Test): Balanced social density for human-robot interactions.

Reasonable Human Density. Humans are grouped by scene area size to balance interaction density. Fig. 3 shows the distribution of human numbers, ensuring meaningful interactions without overcrowding.

Diversity and Scalability. Our datasets include 844 scenes from Social-HM3D and 72 scenes from Social-MP3D, supporting diverse tasks such as SocialNav and social object/image navigation.

B. Experiment Setup

Metrics. Our benchmark metrics focus on task completion and adherence to SocialNav objectives. For task completion, we use Success Rate (Suc.), Success weighted by Path Length (SPL), and Success weighted by Time Length (STL). For social norms, we evaluate Human-Robot Collision Rate (H-coll) and Personal Space Compliance (PSC). Considering the human collision radius (0.3m) and robot size (0.25m), the PSC distance threshold is set to 1.0m.

Baseline Models. We include two classic rule-based methods: A* [48] and ORCA [23]. We compare *Falcon* with Proximity-Aware [18], both are RL-based, and above all methods use depth images as inputs for fair comparison.

Implementation Details. RL agents are trained using DD-PPO [49] with identical hyperparameters. Each algorithm is run three times with different seeds, and results are reported as mean ± standard deviation. Our model is initialized with PointNav weights [46] and fine-tuned for 10 million steps on SocialNav. Training uses 4 Nvidia RTX 3090 GPUs with 8 parallel environments. Models are trained on Social-HM3D and tested on both Social-HM3D and Social-MP3D to evaluate zero-shot generalization.

C. Result Analysis

We conduct experiments to investigate several key aspects:

- Effectiveness of our algorithm against prior methods.
- Impacts of our auxiliary tasks.
- Individual and cooperative effects of the Social Cognition Penalty (SCP) and Spatial-Temporal Precognition Module (SPM).

Table II shows results on Social-HM3D and zero-shot Social-MP3D tests. *Falcon* excels in goal-reaching and social compliance, with strong generalization to unseen envi-

Dataset		Method	Suc. ↑	SPL ↑	STL ↑	PSC ↑	H-Coll ↓
Social-HM3D	Rule-Based	A* [48]	46.14±0.7	46.14±0.7	46.12±0.7	90.56±0.2	53.50±0.9
		ORCA [23]	38.91±0.1	38.91±0.1	38.44±0.1	90.55±0.4	47.52±1.7
	RL	Proximity-Aware [18]	20.11±1.3	18.57±1.9	19.51±1.5	**92.91**±0.5	**33.99**±0.7
		Falcon	**55.15**±0.6	**55.15**±0.7	**54.94**±0.7	89.56±1.4	42.96±1.1
Social-MP3D	Rule-Based	A* [48]	43.85±0.3	43.85±0.3	43.85±0.3	86.74±3.4	57.94±1.5
		ORCA [23]	40.38±0.3	40.38±0.3	39.51±0.2	91.76±0.4	47.16±0.2
	RL	Proximity-Aware [18]	18.45±1.4	17.09±2.8	16.41±1.5	**93.37**±0.9	**32.18**±3.3
		Falcon	**55.05**±0.7	**55.04**±0.6	**54.80**±1.0	90.01±1.2	42.19±0.9

TABLE II: Performance Evaluation of SocialNav Tasks for Rule-Based and RL-Based Methods on Social-HM3D (upper group) and Social-MP3D (lower group). Data in the table represents percentages. We **bold** the best results and underline the second best results.

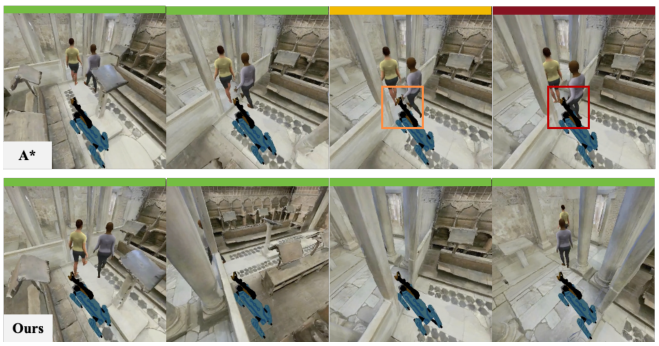

(a) Person Following: A* causes collisions, while our method succeeds.

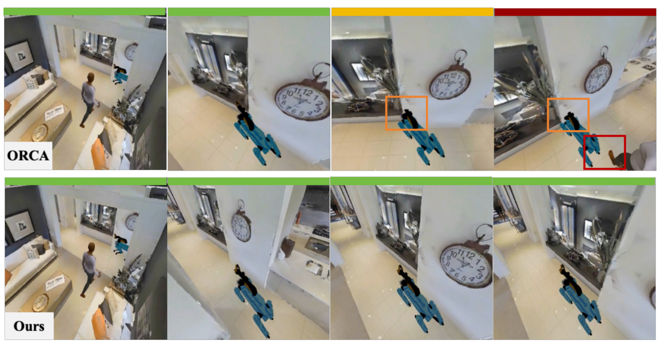

(b) Intersection Encounter: ORCA hits wall, while ours avoids safely.

(c) Frontal Approach: The Proximity-Aware method collides when crossing in front directly, while our method avoids safely by anticipating human path.

Fig. 4: Comparisons of SocialNav Algorithms in Different Encounters: Our method outperforms other algorithms across various encounters. Green indicates safe behaviors, orange indicates risky behaviors (e.g., proximity to humans or collisions with obstacles), and red indicates unsafe behaviors (i.e., collisions with humans).

ronments. Fig. 4 illustrates qualitative comparisons where *Falcon* outperforms others.

Finding 1: Future-aware methods outperform static and situation-aware approaches. Static algorithms like A* fail in dynamic environments (Fig. 4(a)). Situation-aware methods like ORCA and Proximity-Aware react to current states but struggle with delayed responses and collisions (Fig. 4(b), Fig. 4(c)). *Falcon* proactively adjusts to human movements, achieving better performance.

Finding 2: Auxiliary tasks improve performance, with trajectory prediction being most impactful. Table III shows that trajectory forecasting (SPM.Traj) boosts success rates from 40.94% to 54.00%, highlighting its importance.

SPM. Count	SPM. Pos	SPM. Traj	SCP	Suc. ↑	SPL ↑	STL ↑	PSC ↑	H-Coll ↓
PointNav (w/o Aux. Task)				40.94	34.14	11.50	**90.82**	53.54
✓				51.43	51.42	51.16	90.53	46.46
	✓			53.17	53.17	52.95	90.06	44.07
		✓		54.00	53.99	53.92	89.46	43.88
			✓	51.24	51.24	51.08	90.41	48.11
✓	✓	✓		53.63	53.63	53.40	89.33	44.89
✓	✓	✓	✓	**55.15**	**55.15**	**54.94**	89.56	**42.96**

TABLE III: Ablation Study for *Falcon*. The model trained solely with the PointNav algorithm [46] serves as the baseline. SPM.Count, SPM.Pos, and SPM.Traj refer to three auxiliary tasks: Humanoid Count Estimation, Current Position Tracking, and Future Trajectory Forecasting. Data in the table are percentages.

Finding 3: SCP enhances SPM integration, improving performance and training speed. As shown in Table III, SCP significantly improves SPM's performance (55.15% vs. 53.63%).

V. CONCLUSIONS

We introduce a novel SocialNav benchmark with two datasets, Social-HM3D and Social-MP3D, and propose *Falcon*, a future-aware method for social navigation in realistic human-populated scenes. By integrating trajectory prediction into navigation policy, *Falcon* demonstrates superior success rates and collision avoidance. We believe this work will advance research and applications in social navigation.

References

[1] C. Mavrogiannis, F. Baldini, A. Wang, D. Zhao, P. Trautman, A. Steinfeld, and J. Oh, "Core challenges of social robot navigation: A survey," *ACM Transactions on Human-Robot Interaction*, vol. 12, no. 3, pp. 1–39, 2023.

[2] A. Kapoor, S. Swamy, P. Bachiller, and L. J. Manso, "Socnavgym: a reinforcement learning gym for social navigation," in *2023 32nd IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE, 2023, pp. 2010–2017.

[3] W. Wang, L. Mao, R. Wang, and B.-C. Min, "Multi-robot cooperative socially-aware navigation using multi-agent reinforcement learning," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 12 353–12 360.

[4] N. Hirose, D. Shah, K. Stachowicz, A. Sridhar, and S. Levine, "Selfi: Autonomous self-improvement with reinforcement learning for social navigation," *arXiv preprint arXiv:2403.00991*, 2024.

[5] C. Chen, Y. Liu, S. Kreiss, and A. Alahi, "Crowd-robot interaction: Crowd-aware robot navigation with attention-based deep reinforcement learning," in *2019 international conference on robotics and automation (ICRA)*. IEEE, 2019, pp. 6015–6022.

[6] K. Li, Y. Xu, J. Wang, and M. Q.-H. Meng, "Sarl: Deep reinforcement learning based human-aware navigation for mobile robot in indoor environments," in *2019 IEEE International Conference on Robotics and Biomimetics (ROBIO)*. IEEE, 2019, pp. 688–694.

[7] C. Pérez-D'Arpino, C. Liu, P. Goebel, R. Martín-Martín, and S. Savarese, "Robot navigation in constrained pedestrian environments using reinforcement learning," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 1140–1146.

[8] J. Liang, L. Jiang, J. C. Niebles, A. G. Hauptmann, and L. Fei-Fei, "Peeking into the future: Predicting future person activities and locations in videos," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 5725–5734.

[9] H. Nishimura, B. Ivanovic, A. Gaidon, M. Pavone, and M. Schwager, "Risk-sensitive sequential action control with multi-modal human trajectory forecasting for safe crowd-robot interaction," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 11 205–11 212.

[10] 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*, 2021, pp. 15 233–15 242.

[11] J. Liang, L. Jiang, K. Murphy, T. Yu, and A. Hauptmann, "The garden of forking paths: Towards multi-future trajectory prediction," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 10 508–10 518.

[12] J. Liang, L. Jiang, and A. Hauptmann, "Simaug: Learning robust representations from simulation for trajectory prediction," in *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIII 16*. Springer, 2020, pp. 275–292.

[13] H. Zhao, J. Gao, T. Lan, C. Sun, B. Sapp, B. Varadarajan, Y. Shen, Y. Shen, Y. Chai, C. Schmid, *et al.*, "Tnt: Target-driven trajectory prediction," in *Conference on Robot Learning*. PMLR, 2021, pp. 895–904.

[14] A. Francis, C. Pérez-d'Arpino, C. Li, F. Xia, A. Alahi, R. Alami, A. Bera, A. Biswas, J. Biswas, R. Chandra, *et al.*, "Principles and guidelines for evaluating social robot navigation algorithms," *arXiv preprint arXiv:2306.16740*, 2023.

[15] S. Liu, P. Chang, W. Liang, N. Chakraborty, and K. Driggs-Campbell, "Decentralized structural-rnn for robot crowd navigation with deep reinforcement learning," in *2021 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2021, pp. 3517–3524.

[16] P. Kothari, S. Kreiss, and A. Alahi, "Human trajectory forecasting in crowds: A deep learning perspective," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 7, pp. 7386–7400, 2022.

[17] S. Holk, D. Marta, and I. Leite, "Polite: Preferences combined with highlights in reinforcement learning," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 2288–2295.

[18] E. Cancelli, T. Campari, L. Serafini, A. X. Chang, and L. Ballan, "Exploiting proximity-aware tasks for embodied social navigation," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 10 957–10 967.

[19] F. Xia, W. B. Shen, C. Li, P. Kasimbeg, M. E. Tchapmi, A. Toshev, R. Martín-Martín, and S. Savarese, "Interactive gibson benchmark: A benchmark for interactive navigation in cluttered environments," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 713–720, 2020.

[20] X. Puig, E. Undersander, A. Szot, M. D. Cote, T.-Y. Yang, R. Partsey, R. Desai, A. W. Clegg, M. Hlavac, S. Y. Min, *et al.*, "Habitat 3.0: A co-habitat for humans, avatars and robots," *arXiv preprint arXiv:2310.13724*, 2023.

[21] 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*, p. 02783649241230562, 2024.

[22] R. Möller, A. Furnari, S. Battiato, A. Härmä, and G. M. Farinella, "A survey on human-aware robot navigation," *Robotics and Autonomous Systems*, vol. 145, p. 103837, 2021.

[23] J. Van Den Berg, S. J. Guy, M. Lin, and D. Manocha, "Reciprocal n-body collision avoidance," in *Robotics Research: The 14th International Symposium ISRR*. Springer, 2011, pp. 3–19.

[24] J. Van den Berg, M. Lin, and D. Manocha, "Reciprocal velocity obstacles for real-time multi-agent navigation," 2008.

[25] P. Long, W. Liu, and J. Pan, "Deep-learned collision avoidance policy for distributed multiagent navigation," *IEEE Robotics and Automation Letters*, vol. 2, no. 2, pp. 656–663, 2017.

[26] G. S. Aoude, B. D. Luders, J. M. Joseph, N. Roy, and J. P. How, "Probabilistically safe motion planning to avoid dynamic obstacles with uncertain motion patterns," *Autonomous Robots*, vol. 35, no. 1, pp. 51–76, 2013.

[27] J. Guzzi, A. Giusti, L. M. Gambardella, G. Theraulaz, and G. A. Di Caro, "Human-friendly robot navigation in dynamic environments," in *IEEE international conference on robotics and automation*. IEEE, 2013, pp. 423–430.

[28] G. Ferrer, A. Garrell, and A. Sanfeliu, "Social-aware robot navigation in urban environments," in *Proc. of thr European Conference on Mobile Robots*, 2013.

[29] Y. F. Chen, M. Everett, M. Liu, and J. P. How, "Socially aware motion planning with deep reinforcement learning," 2017.

[30] Y. Lu, X. Ruan, and J. Huang, "Deep reinforcement learning based on social spatial–temporal graph convolution network for crowd navigation," *Machines*, vol. 10, no. 8, p. 703, 2022.

[31] C. Chen, Y. Liu, S. Kreiss, and A. Alahi, "Crowd-robot interaction: Crowd-aware robot navigation with attention-based deep reinforcement learning," 2019.

[32] 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*, vol. 39, no. 8, pp. 895–935, 2020.

[33] R. Huang, H. Xue, M. Pagnucco, F. Salim, and Y. Song, "Multimodal trajectory prediction: A survey," *arXiv preprint arXiv:2302.10463*, 2023.

[34] A. Elnagar, "Prediction of moving objects in dynamic environments using kalman filters," in *Proceedings 2001 IEEE International Symposium on Computational Intelligence in Robotics and Automation (Cat. No. 01EX515)*. IEEE, 2001, pp. 414–419.

[35] S. Zernetsch, S. Kohnen, M. Goldhammer, K. Doll, and B. Sick, "Trajectory prediction of cyclists using a physical model and an artificial neural network," in *2016 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2016, pp. 833–838.

[36] P. Coscia, F. Castaldo, F. A. Palmieri, A. Alahi, S. Savarese, and L. Ballan, "Long-term path prediction in urban scenarios using circular distributions," *Image and Vision Computing*, vol. 69, pp. 81–91, 2018.

[37] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, and S. Savarese, "Social lstm: Human trajectory prediction in crowded spaces," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 961–971.

[38] T. P. Kucner, M. Magnusson, E. Schaffernicht, V. H. Bennetts, and A. J. Lilienthal, "Enabling flow awareness for mobile robots in partially observable environments," *IEEE Robotics and Automation Letters*, vol. 2, no. 2, pp. 1093–1100, 2017.

[39] A. Vemula, K. Muelling, and J. Oh, "Modeling cooperative navigation in dense human crowds," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 1685–1692.

[40] D. Vasquez, "Novel planning-based algorithms for human motion prediction," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2016, pp. 3317–3322.

[41] A. Rudenko, L. Palmieri, and K. O. Arras, "Predictive planning for a mobile robot in human environments," in *Proc. of the IEEE Int.*

Conf. on Robotics and Automation (ICRA), Works. on AI Planning and Robotics, 2017.

[42] E. Rehder, F. Wirth, M. Lauer, and C. Stiller, "Pedestrian prediction by planning using deep neural networks," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 5903–5908.

[43] C. Li, J. Jang, F. Xia, R. Martín-Martín, C. D'Arpino, A. Toshev, A. Francis, E. Lee, and S. Savarese, "igibson challenge 2021," Online, 2021, [Online; accessed: 2024-08-23]. [Online]. Available: http://svl.stanford.edu/igibson/challenge.html

[44] V. Makoviychuk, L. Wawrzyniak, Y. Guo, M. Lu, K. Storey, M. Macklin, D. Hoeller, N. Rudin, A. Allshire, A. Handa, *et al.*, "Isaac gym: High performance gpu-based physics simulation for robot learning," *arXiv preprint arXiv:2108.10470*, 2021.

[45] A. D. Vuong, T. T. Nguyen, M. N. VU, B. Huang, D. Nguyen, H. T. T. Binh, T. Vo, and A. Nguyen, "Habicrowd: A high performance simulator for crowd-aware visual navigation," *arXiv preprint arXiv:2306.11377*, 2023.

[46] S. K. Ramakrishnan, A. Gokaslan, E. Wijmans, O. Maksymets, A. Clegg, J. Turner, E. Undersander, W. Galuba, A. Westbury, A. X. Chang, *et al.*, "Habitat-matterport 3d dataset (hm3d): 1000 large-scale 3d environments for embodied ai," *arXiv preprint arXiv:2109.08238*, 2021.

[47] A. Chang, A. Dai, T. Funkhouser, M. Halber, M. Niessner, M. Savva, S. Song, A. Zeng, and Y. Zhang, "Matterport3d: Learning from rgb-d data in indoor environments," *arXiv preprint arXiv:1709.06158*, 2017.

[48] P. E. Hart, N. J. Nilsson, and B. Raphael, "A formal basis for the heuristic determination of minimum cost paths," *IEEE transactions on Systems Science and Cybernetics*, vol. 4, no. 2, pp. 100–107, 1968.

[49] E. Wijmans, A. Kadian, A. Morcos, S. Lee, I. Essa, D. Parikh, M. Savva, and D. Batra, "Dd-ppo: Learning near-perfect pointgoal navigators from 2.5 billion frames," *arXiv preprint arXiv:1911.00357*, 2019.