

Learning to Socially Navigate in Pedestrian-rich Environments with Interaction Capacity

Quecheng Qiu¹, Shunyi Yao², Jing Wang¹, Jun Ma¹, Guangda Chen², and Jianmin Ji^{2,*}

Abstract—Existing navigation policies for autonomous robots tend to focus on collision avoidance while ignoring human-robot interactions in social life. For instance, robots can pass along the corridor safer and easier if pedestrians notice them. Sounds have been considered as an efficient way to attract the attention of pedestrians, which can alleviate the freezing robot problem. In this work, we present a new deep reinforcement learning (DRL) based social navigation approach for autonomous robots to move in pedestrian-rich environments with interaction capacity. Most existing DRL based methods intend to train a general policy that outputs both navigation actions, i.e., expected robot’s linear and angular velocities, and interaction actions, i.e., the beep action, in the context of reinforcement learning. Different from these methods, we intend to train the policy via both supervised learning and reinforcement learning. In specific, we first train an interaction policy in the context of supervised learning, which provides a better understanding of the social situation, then we use this interaction policy to train the navigation policy via multiple reinforcement learning algorithms. We evaluate our approach in various simulation environments and compare it to other methods. The experimental results show that our approach outperforms others in terms of the success rate. We also deploy the trained policy on a real-world robot, which shows a nice performance in crowded environments.

I. INTRODUCTION

Mobile robots have been widely used in various real-world applications, where the capability of autonomous navigation is critical. Most of these applications require the robot to navigate in pedestrian-rich environments, like malls, hotels, and hospitals. It is challenging for the robot to shuttle freely in the crowd while following social conventions.

Most existing work for robot navigation focus on collision avoidance while ignoring human-robot interactions in social life. Some of them navigate the robot by first analyzing pedestrians’ behaviors and predicting their trajectories [1]–[4]. However, it is hard to precisely predicate pedestrians’ trajectories and is like to suffer from the freezing robot problem in dense crowds. On the other hand, deep reinforcement learning (DRL) based methods have been applied

to collision avoidance with promising results [5]–[9], where the reward functions need to be specified to describe how the robot ought to move. Most of them ignore human-robot interactions, regard pedestrians as moving obstacles, and plan the avoiding-path over the range of pedestrians’ moving area. However, it would be much more efficient and safer if the robot were noticed by pedestrians, which brings new challenges for the robot on when to perform interaction actions and how to navigate in this human-aware fashion.

Interaction via sounds has been considered as an efficient way to attract the attention of pedestrians. Nishimura et al. [10] design a DRL framework to learn general policies that output both navigation actions, i.e., expected robot’s linear and angular velocities, and interaction actions, i.e., the beep action. However, there are several drawbacks for training these general policies in the context of DRL: 1) It is difficult to design a specific reward function that encourages the robot to perform interaction actions properly in various social environments, like a narrow corridor with moving pedestrians or a hall full of standing people. 2) Moreover, the social conventions followed by people are hard to be specified by reward functions. 3) The general policy for both navigation actions and interaction actions results in a much larger policy space to be explored, and the nature of the sparse reward for navigation requires a long-term memory to store the succeeding states and actions after an interaction action during the training.

In this paper, we address these problems by training the policy via both DRL and supervised learning (SL), as shown in Fig. 1. In specific, we first train an interaction policy in the context of SL. Note that, it is easier for humans to label the situations when an interaction action should be performed following social conventions, other than specifying such a reward function. Then a proper interaction policy that is suitable for the certain social environment can be learned via SL from the labeled interaction data. Later, we use this learned interaction policy to train the navigation policy via DRL. Multiple DRL methods have been applied for robot navigation with promising results. We can slightly adjust these methods to train the navigation policy given the learned interaction policy. We implement our approach based on multiple DRL algorithms, evaluate the approach in various simulation environments, and compare it to other methods. The experimental results show that our approach outperforms others in terms of the success rate. We also deploy the trained policy on a real-world robot, which shows a nice performance in crowded environments.

Our main contributions are summarized as follows:

The work is partially supported by the National Key Research and Development Program of China (No. 2018AAA0100500), CAAL-Huawei MindSpore Open Fund, Anhui Provincial Development and Reform Commission 2020 and 2021 Projects, Key-Area Research and Development Program of Guangdong Province 2020B0909050001, Shenzhen Yijiahe Technology R&D Co., Ltd., and Huawei Cloud Computing Technologies Co., Ltd.

* The corresponding author.

¹ School of Data Science, University of Science and Technology of China, Hefei, 230026, China {qiuqc, wj980828, markjun}@mail.ustc.edu.cn.

² School of Computer Science and Technology, USTC, Hefei, 230026, China {ustcysy, cgdsss}@mail.ustc.edu.cn, jianmin@ustc.edu.cn.

- We propose a learning approach for social navigation with interaction capacity in pedestrian-rich environments. We first train a Supervised Learning Interaction (SLI) policy, then we use this learned interaction policy to train the navigation policy via DRL.
- We implement our approach based on multiple DRL algorithms, which implies that the approach can be used to incorporate other DRL based navigation methods with interaction policies.
- We evaluate the approach in various simulation environments and compare it to other methods. The experimental results show that our approach outperforms others in terms of the success rate. We also deploy the trained policy on a real-world robot, which shows a nice performance in crowded environments.

II. RELATED WORK

A. Crowd-avoidance Robot Navigation.

It is challenging for robots to navigate in pedestrian-rich environments since the prediction trajectories of pedestrians are hardly to be precise. Prior work address the problem by introducing various reciprocal methods, like SFM [3], RVO [1], ERVO [11] and ORCA [2], which have achieved good results both in both simulation and real-world testing environments. However, these methods rely on certain assumptions about patterns of human behavior, which can hardly handle the scenarios with dynamic groups that exhibit different spatial behavior in pedestrian-rich environments. Other methods, like Fuzzy control [12], are used to predict the movement of nearby human [13].

On the other hand, DRL based methods have been applied to crowd avoidance navigation tasks with promising results. These DRL methods can generally be divided into two categories, i.e., value-based methods and policy-based methods. Value-based methods learn policy based on value function while policy-based methods based on gradient boost. For both of them, accurately estimating the state-value function is vital for the navigation performance, which is highly relied on the presentation of agents' observation. For value-based methods, Chen et al. [6] propose an agent-level method called SA-CADRL and Chen et al. [8] employ Graph Convolutional Networks for a better representation of the crowd as well. Both of them have made promising achievements in social navigation tasks. However, they are lack of support for continuous actions and are not good at searching for stochastic optimal policies [14]. As for policy-based methods, Fan et al. [15] propose a hybrid-RL policy for multi-agent collision avoidance without pedestrian information, which also works in dense crowds. Liu et al. [7] separate pedestrians from static obstacles based on the angular map. Yao et al. [16] use both the egocentric sensor map and the pedestrian map as the real-time observation, which performs well in pedestrian-rich environments without human-robot interactions.

Imitation learning methods [17] aim to learn a policy as close to the expert policy as better, which requires the collection of experts' demonstrations. Tai et al. [18] applies GAIL [19] to navigation tasks with raw depth inputs, which

combines the idea of GAN [20] and imitation learning and performs well in realistic scenarios.

B. Human-Robot Interaction in Robot Navigation.

More and more robots are running in dynamic environments with humans, it's important for autonomous robots to interact with humans following the social conventions. Kruse et al. [21] summarize that, comfort, naturalness, and sociability are three key requirements of social navigation.

Prior work usually use model-based methods to simulate the interaction between robots and crowds. Some early work regard the crowd as static obstacles [22]–[24], and avoid collision by establishing a virtual force field. Later, the Social Force model [25]–[27] is proposed. In this model, social interactions are simulated by controlling the repulsive forces of different types (human-to-human, human-to-robot), where the force parameters usually need to be adjusted separately, which limits its generalization to many applications. ORCA [2] propose a method for reciprocal n -body collision avoidance under the assumption that the protocol used by each robot is the same. Furthermore, ERVO [11] proposes a new emotional contagion model on the basis of RVO [1].

In contrast, rather than formulating strategies, learning-based methods have received widespread attention to guide the robot to navigate by simulating crowd behavior. Kim and Pineau [28] use inverse reinforcement learning to learn a cost function through human demonstrations. In order to deal with various neighbors, Everett et al. [29] use the LSTM model and process each neighbor in the reverse order of the distances from the robot. However, this model cannot fully represent all the interactions between the crowd and the robot. SARL [30] designs a social attentive pooling module to encode crowd cooperative behaviors, where the human-robot and human-human interaction are jointly modeled, which makes up for the deficiencies in [29]. However, all of the above methods may fall into stuck or time out when the crowd is dense. To address the freezing robot problem, Dugas et al. [31] propose a multi-model socially navigation behaviors which allows robots to speak and nudge. Nishimura et al. [10] design the safety-efficiency trade-off reward based on Sequential Social Dilemmas (SSD) [32]. Both of them propose that instead of getting into the freezing robot problem, it is better to encourage robots to interact with pedestrians appropriately.

Following above ideas, we address the social navigation problem with interaction capacity by training the policy via both DRL and SL.

III. APPROACH

It is neither safe nor efficient for a robot to take the initiative to avoid collisions with every pedestrian in a pedestrian-rich environment. On the other hand, with interaction capacity, the robot can interact with humans when it were going to be stuck, which can alleviate the freezing robot problem and improve efficiency.

In this section, we propose a new learning approach for social navigation with interaction capacity. We first train an

interaction policy via SL, then we use this learned interaction policy to train the navigation policy via DRL. As shown in Fig. 1, the navigation policy network is composed of two modules, i.e.,

- planning module, serves as a collision avoidance component in a typical path planning system for a mobile robot, which observes the environment via its surrounding sensor and pedestrian information and outputs expected linear and angular velocities to navigate the robot,
- interaction module, specifies an interaction policy for robot to decide when to interact with humans, i.e., beep or speak, given its surrounding pedestrian information and its next control commands.




Fig. 1. The architecture of the crowd navigation policy network.

A. Planning Module.

The input for the planning module is composed of three parts, i.e., the sensor information, the pedestrian information, and the relative goal for the robot. In particular, the sensor information is specified by an egocentric local grid map of the robot, which represents the environmental information around the robot, including its shape and observable appearances of obstacles. The local grid map is constructed from a costmap that is generated by outputs of a 2D laser. The pedestrian information consists of three channels of pedestrian maps, which indicate the location and speed of pedestrians around the robot. The relative goal specifies the target position and orientation of the robot.

In this paper, we implement the approach on a differential drive robot that follows desired speed commands. Then the output for the planning module consists of a linear velocity v and an angular velocity ω . We implement both discrete and continuous actions for using multiple DRL methods. In specific, for discrete actions, we set a linear velocity $v \in \{0.0, 1.0\}$ and an angular velocity $\omega \in \{-0.8, -0.4, 0.0, 0.4, 0.8\}$. For continuous actions, we set $v \in [0, 0.6]$ and $\omega \in [-0.9, 0.9]$. Both discrete and continuous actions can be directly performed by the differential robot in our experiments. Note that, $v \geq 0$, i.e., moving backward is not allowed, due to the lack of rear sensors. As shown in Fig. 1, for discrete actions, the network outputs a 10-dimensional vector from a softmax layer to choose the pair

of linear and angular velocities. For continuous actions, the network outputs the mean of linear (resp. angular) velocity sampled from a Gaussian distribution.

Notice that, the learned interaction policy does not directly change the inputs of the planning module. It affects the training of the planning module by affecting behaviors of surrounding pedestrians in the training simulation environments for the planning module in the context of DRL. Then we can only ask the navigation policy to minimize the arriving time of the robot without collision. When the policy-based DRL algorithm, i.e., PPO [33], is applied in the planning module, we follow the reward specified in [16].

The network for the planning module would be trained in our customized simulator¹. In our experiments, both the valued-based DRL algorithm, i.e., V-learning [5], and the policy-based DRL algorithm, i.e., PPO, would be implemented and evaluated in the simulator for the discrete and continuous actions, respectively.

B. Interaction Module.

When to beep or alarm is vital for the interaction policy in social robot navigation, which requires the robot to better understanding the social situation. A easy policy P_{beep}^1 is to perform the beep action A_{beep} , when d_{min} (the distance from the robot to the nearest pedestrian) is smaller than a fixed distance d_θ , i.e.,

$$P_{beep}^1 = \begin{cases} A_{beep}, & \text{if } d_{min} < d_\theta, \\ Do\ nothing, & \text{otherwise.} \end{cases} \quad (1)$$

We can construct an improved policy P_{beep}^2 by further considering the speed and orientation of pedestrians, i.e.,

$$P_{beep}^2 = \begin{cases} A_{beep} & \text{if } \forall p. \|\mathbf{P}_p\| < d_\theta \wedge \|\mathbf{V}_p\| > v_\theta \\ & \wedge \mathbf{P}_p \cdot \mathbf{V}_p > 0, \\ Do\ nothing, & \text{otherwise.} \end{cases} \quad (2)$$

Where \mathbf{P}_p is the relative position between the robot and the pedestrian p , \mathbf{V}_p denotes the velocity of the pedestrian, and v_θ is the velocity threshold of the pedestrian.

Above simple interaction policies require predefined fixed thresholds, which may cause unnatural behaviors and are not robust for various dense environments. Nishimura and Yonetani introduce L2B [10] that designs a delicate reward function to let the robot learn to beep in the context of DRL.

To better understand the social situation, we propose to learn the interaction policy from the interaction data labeled by human, where the social conventions could be implicitly learned via SL. Moreover, the interaction policy should not only perform the beep action when it has already got stuck, but also try its best to prevent the robot from freezing. In this sense, the next control command of the robot, i.e., the expected linear and angular velocities, needs to be considered as an input of the network for the interaction policy and should be considered when the data labeling as well.

¹https://github.com/DRL-Navigation/img_env.

As shown in Fig. 1, after training the network by the labeled interaction data, we can integrate the learned interaction policy to the network for the navigation policy and train the integrated navigation policy in the context of DRL. Notice that, in the training process of DRL, the robot is able to observe multiple steps of following states after the beep action. Then the integrated navigation policy would be able to learn to cooperate with the interaction policy to prevent it from freezing. We will evaluate such performances in the experiments.

As shown in Fig. 1, the convolutional neural network for binary classification is used in the interaction module for SL. The network consists of three convolutional layers to extract features hidden in the pedestrian map. Then, the refined features are merged with the robot’s control command from the planning module and passed to a fully-connected layer, which outputs a 2-dimensional vector from a softmax layer to choose whether beep or not.

C. Simulation Environments.

We train both the planning module and the interaction module in simulation environments generated by a customized 2D simulator based on OpenCV. Following the discussion in [9], [16], training the navigation policy in two kinds of scenarios, i.e., random and circular scenarios as shown in Fig. 2, can lead to a robust policy that not only performs well in various simulation environments but also in the real world.

In this paper, we train both the planning and interaction policies in these two scenarios. In particular, a random scenario contains one robot, eight pedestrians, and four static obstacles, where their initial positions and ending positions (for the robot and pedestrians) would be placed randomly. Meanwhile, a circular scenario contains one robot and eight pedestrians, where all of them would be placed randomly in a circle with a random radius.

Following L2B², all pedestrians in our simulator are driven by Emotional Reciprocal Velocity Obstacles (ERVO) [11], an emotional version of RVO which allows pedestrians to react with the beep sound naturally, where pedestrians’ max speed is set to be 1 m/s.

IV. EXPERIMENTS

In this section, we evaluate our approach for social navigation with interaction capacity in both the simulation and the real world. We first specify details of our implementation on both the planning module and the interaction module. Then we compare our approach with others, including the ones with simple interaction policies P_{beep}^1 and P_{beep}^2 , and L2B. We demonstrate the robustness of the approach by implementing it on different DRL algorithms. We also deploy the trained model to a differential drive robot and test its navigation performance in the real world. Both qualitative and quantitative experiments show that our approach performs well in pedestrian-rich environments.

²<https://github.com/denkiwakame/Python-ERVO>.




Fig. 2. Two scenarios for the training, where the small black circle and light blue spots behind it denote the running robot and its trajectories, green spots denote the legs’ trajectories of pedestrians, other black boxes and circles denote static obstacles, and the blue circle line around the robot denotes the warning zone for the beep action w.r.t. nearby pedestrians.

TABLE I
PARAMETERS FOR SL AND DRL

SL hyper-parameters	Value	DRL hyper-parameters	Value
dataset size	10000	learning rate for policy	5×10^{-5}
batch size	1024	learning rate for value	1×10^{-3}
image size	48×48	discount factor (γ)	0.99
pedestrian radius	0.3	replay buffer size	2048
learning rate of Adam	1×10^{-4}	image size	48×48
training epoch	10000	maximum episode length	200

A. Supervised Learning and Reinforcement Learning Setup.

We start with the training of the interaction module via SL. The labeled training data contains 10,000 pairs of (x, y) , where x is a triple $\langle M_P, v, \omega \rangle$, M_P denotes the pedestrian information which is specified by three channels of the grayscale images to represent the location and speed of pedestrians around the robot, v and ω denote the linear and angular velocities of the robot, and y is the label for whether beep or not. These training data are collected from various simulation environments in the two scenarios, i.e., random scenario and circular scenario. We label these data manually. For DRL, we use PPO as the policy-based method, and V-learning as the value-based method, the parameters of V-learning are all the same as in SARL³. Parameters for SL and PPO are listed in Table I.

In our experiments, both SL and DRL are trained in an i9-9900k CPU and an NVIDIA Titan RTX GPU.

B. Experiments on Simulation Environments.

1) *Comparison with simple interaction policies:* We compare the performance of our approach, i.e., SLI, with the ones that replacing the learning interaction policy by simple policies P_{beep}^1 and P_{beep}^2 respectively, i.e., $FD(d_\theta)$ and $FDV(d_\theta, v_\theta)$, where $FD(d_\theta)$ denotes the social navigation policy that the robot would beep when a pedestrian gets close to a distance less than d_θ as defined in Eq. (1), and $FDV(d_\theta, v_\theta)$ considers both the distance d_θ and the velocity v_θ . Note that, the velocity threshold only considers pedestrians who move towards the robot as specified in Eq. (2). We set $d_\theta = 1.0m$ as the distance threshold and choose $v_\theta \in \{0.3, 0.5, 0.7\} m/s$

³<https://github.com/vita-epfl/CrowdNav>.

as possible velocity thresholds. In addition, we use Base to denote the baseline approach that omits the interaction module in SLI. Here we apply PPO in the planning module for the above methods, which outputs continuous actions of the robot.

Table II shows the average results of these methods in 500 testing environments of each scenario, where ‘Success’ denotes the ratio of arriving at the goal safely, ‘PedColl’ denotes the ratio of collisions with pedestrians, and ‘Beep’ denotes the ratio of beep times in all steps.

TABLE II
PERFORMANCE OF DIFFERENT INTERACTION METHODS

scenarios	Methods	Success	PedColl	Beep
Random	Base	0.668	0.330	-
	FD(1.0)	0.956	0.010	0.624
	FDV(1.0, 0.3)	0.912	0.056	0.120
	FDV(1.0, 0.5)	0.866	0.100	0.077
	FDV(1.0, 0.7)	0.804	0.138	0.034
	SLI	0.884	0.062	0.097
Circular	Base	0.558	0.442	-
	FD(1.0)	0.778	0.222	0.823
	FDV(1.0, 0.3)	0.830	0.170	0.256
	FDV(1.0, 0.5)	0.732	0.268	0.254
	FDV(1.0, 0.7)	0.766	0.222	0.179
	SLI	0.864	0.136	0.158

Table II shows interaction policies can greatly improve the performance of robot navigation in both random and circular scenarios. Moreover, our approach, SLI, performs well in environments with high-frequent interactions, like environments in the circular scenario. Note that, SLI leads to a much lower beep rate than the fixed-threshold methods, i.e., FD and FDV, while keeping a high success rate. Fig. 3 demonstrates the differences of the beep policy between FDV(1.0, 0.5) and SLI. When the robot is surrendered by a crowd, FDV beeps frequently by the fixed-threshold policy P_{beep}^2 , while SLI achieves the similar performance with only few beeps.

For environments in the random scenario, SLI still leads to an acceptable performance compared with others. Meanwhile, proper values of d_θ and v_θ need to be manually adjusted to achieve better performance. These experimental results show that a proper interaction policy can greatly improve the performance of robot navigation in pedestrian-rich environments with very few interactions.

2) *Comparison with L2B-SARL*: Nishimura and Yonetani introduce L2B that designs a delicate reward function to let the robot learn to beep in the context of DRL. They propose L2B-SARL as a combination of L2B and SARL, i.e., a self-attention mechanism.

We implement L2B-SARL and train it on our simulator. For a fair comparison, we also implement V-SLI, which is extended from L2B-SARL by replacing the policy for the beep action with the interaction policy trained by SL in our approach. Note that, both approaches output discrete actions of the robot.

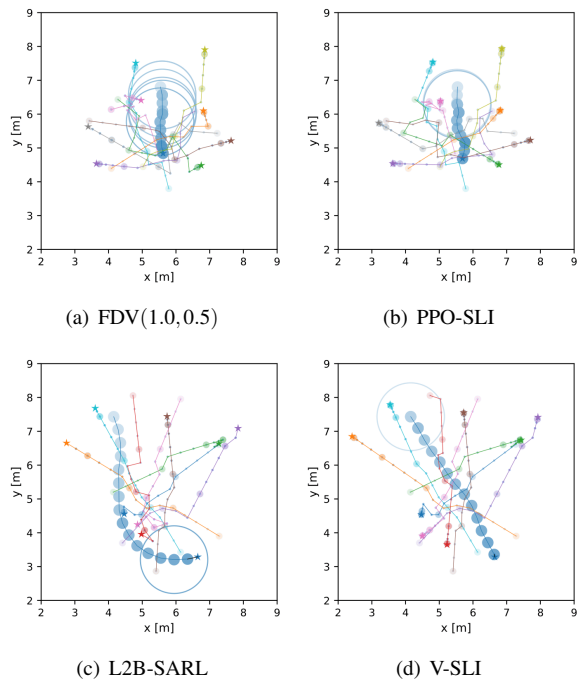


Fig. 3. The illustration of trajectories generated by various social navigation approaches in the circular scenario with one robot and eight pedestrians. The gradient blue spots denote the trajectory of the running robot, the blue circle around the robot denotes the beep action, and lines with other colors denote different pedestrians, where stars denote their target points.

For a fair comparison, we remove static obstacles in environments of the random scenario here, as L2B-SARL does not support static obstacles. In the experiments, the network parameters were pre-trained via imitation learning, i.e., 6k episodes by ORCA firstly, which due to the fact that the reward function of L2B does not encourage the robot to move towards the goal.

Both quantitative and qualitative results are presented in the following, which show a clear superiority of our method.

Table III lists the numerical results of the comparison experiments under different pedestrian numbers N in 500 testing environments of each scenario. The results show that V-SLI performs better than L2B-SARL in all tests in the sense of success rate and collision rate, while it leads to a slightly higher frequency of the beep action. With the increase of the number of pedestrians, the performance of L2B-SARL decreases rapidly, while V-SLI still performs well. Moreover, due to the existence of saddle points in state space about the value function, L2B-SARL has a higher stuck rate than ours, since their action space is twice as large as ours.

Fig. 3 illustrates the difference of interaction policies between L2B-SARL and V-SLI. At the beginning of the trajectory, there is a pedestrian near the robot, L2B-SARL chooses to change the robot’s direction to avoid the collision, which leads the robot to take a big turn and meet more pedestrians. Meanwhile, V-SLI predicts the collision and decides to beep at once, then efficiently moves straight forward to the goal.

TABLE III
PERFORMANCE OF L2B-SARL AND V-SLI

scenarios	Methods	N	Success	PedColl	Beep	
Random	L2B-SARL	4	0.824	0.168	0.052	
	V-SLI	4	0.920	0.080	0.065	
	L2B-SARL	6	0.686	0.198	0.046	
	V-SLI	6	0.882	0.108	0.098	
	L2B-SARL	8	0.508	0.198	0.048	
	V-SLI	8	0.770	0.144	0.094	
	L2B-SARL	10	0.678	0.286	0.051	
	V-SLI	10	0.840	0.100	0.106	
	Circular	L2B-SARL	4	0.694	0.296	0.044
		V-SLI	4	0.808	0.190	0.081
L2B-SARL		6	0.546	0.442	0.047	
V-SLI		6	0.656	0.342	0.109	
L2B-SARL		8	0.520	0.270	0.047	
V-SLI		8	0.612	0.378	0.179	
L2B-SARL		10	0.514	0.278	0.046	
V-SLI		10	0.562	0.360	0.103	

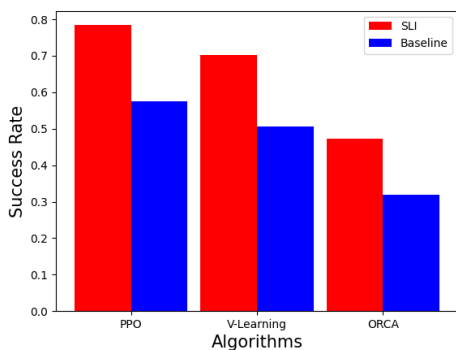


Fig. 4. Success rate of PPO-SLI, V-SLI and ORCA-SLI with respect to their baselines. The average results are tested in 1000 episodes which consisted of 500 environments of the random scenario and 500 environments of the circle scenario, with 10 pedestrians.

3) *Generalizable experiments*: We implement the SL interaction module on both the value-based method, i.e., V-learning and the policy-based method, i.e., PPO, and we also implement the SL interaction module on a conventional model-based method, i.e., ORCA. The performance of these approaches, i.e., *success rate*, is shown in Fig. 4, where the baselines denote the original methods without the SL interaction module.

Fig. 4 shows that the SL interaction module can be easily adapted to multiple navigation methods and improve their performance. The common promotion for baselines with interaction module demonstrates that our approach is generalizable and advantageous.

C. Experiments on the Real Robot.

As shown in Fig. 5, we use Kobuki base TurtleBot 2 as the real-world robot, which is equipped with a Hokuyo UTM-30LX Scanning Laser Rangefinder as the 2D laser sensor, a RealSense D455 depth camera as the pedestrian detection and tracking sensor, as well as a Meeteasy Mvoice 1000 as the audio output. In addition, we choose ThinkPad P15v as




Fig. 5. The robot in our experiments, where the 3-channel pedestrian map is generated from the depth camera, the sensor map is from the laser sensor, the laptop is for inferring the deep neural network (left). A pedestrian-rich environment (right) and the beeping robot (bottom-right) are illustrated.




Fig. 6. The illustration of the real-world robot, which tracks persons through combining yolo and spencer people tracking (left). The trajectory of the robot (blue cycle) in a crowded environment (right) is illustrated, where the red circle denotes its beep action.

the computation device (Intel core i7-11800H CPU, 8cores, 16threads, 4.6GHZ, Nvidia-T600 4GB GPU, 2.1kg), which allows the real robot to run the control model at 63 HZ. We deploy the trained policy by PPO-SLI to the robot since it outperforms others including V-SLI. We use yolo-V3 [34] to detect pedestrians and spencer people tracking [35] to track pedestrians, as shown in Fig. 6.

We test our method in a crowded corridor environment, as shown in Fig. 5, the robot is carrying out a navigation task while pedestrians walking in the vicinity. Note that, some humans may ignore the robot, i.e., is using the mobile phone, which increases the difficulty of navigation task. Fig. 6 illustrates the trajectory of the robot in the environment, which demonstrates that our approach allows the robot reach the goal efficiently and beep in a proper way while people are about to surround it. Furthermore, when obscured by pedestrians, proper interactions to remind pedestrians to move away are beneficial for the robot to restore its localization.

V. CONCLUSION

We propose a learning approach for robot social navigation with interaction capacity in pedestrian-rich environments, where the interaction policy, SLI, is first trained via SL, then it incorporates with a DRL based navigation method to generate the social navigation policy. We evaluate the approach on both simulation environments and a real-world robot. Both quantitative and qualitative experimental results show that, the SLI policy can be easily adapted to multiple navigation methods and improve their performance, which allows the robot to navigate through crowds efficiently.

REFERENCES

- [1] J. Van den Berg, M. Lin, and D. Manocha, "Reciprocal velocity obstacles for real-time multi-agent navigation," in *Proceedings of ICRA-2008*. IEEE, 2008, pp. 1928–1935.
- [2] J. Van Den Berg, S. J. Guy, M. Lin, and D. Manocha, "Reciprocal n-body collision avoidance," in *Robotics research*. Springer, 2011, pp. 3–19.
- [3] D. Helbing and P. Molnar, "Social force model for pedestrian dynamics," *Physical review E*, vol. 51, no. 5, p. 4282, 1995.
- [4] M. Phillips and M. Likhachev, "Sipp: Safe interval path planning for dynamic environments," in *2011 IEEE International Conference on Robotics and Automation*. IEEE, 2011, pp. 5628–5635.
- [5] Y. F. Chen, M. Liu, M. Everett, and J. P. How, "Decentralized non-communicating multiagent collision avoidance with deep reinforcement learning," in *2017 IEEE international conference on robotics and automation (ICRA)*, 2017, pp. 285–292.
- [6] Y. F. Chen, M. Everett, M. Liu, and J. P. How, "Socially aware motion planning with deep reinforcement learning," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017, pp. 1343–1350.
- [7] L. Liu, D. Dugas, G. Cesari, R. Siegart, and R. Dubé, "Robot navigation in crowded environments using deep reinforcement learning," in *Proceedings of IROS-2020*, 2020, pp. 5671–5677.
- [8] Y. Chen, C. Liu, B. E. Shi, and M. Liu, "Robot navigation in crowds by graph convolutional networks with attention learned from human gaze," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 2754–2761, 2020.
- [9] G. Chen, S. Yao, J. Ma, L. Pan, Y. Chen, P. Xu, J. Ji, and X. Chen, "Distributed non-communicating multi-robot collision avoidance via map-based deep reinforcement learning," *Sensors*, vol. 20, no. 17, p. 4836, 2020.
- [10] M. Nishimura and R. Yonetani, "L2b: Learning to balance the safety-efficiency trade-off in interactive crowd-aware robot navigation," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 11 004–11 010.
- [11] M. Xu, X. Xie, P. Lv, J. Niu, H. Wang, C. Li, R. Zhu, Z. Deng, and B. Zhou, "Crowd behavior simulation with emotional contagion in unexpected multihazard situations," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 2019.
- [12] L. Zadeh, "Fuzzy sets," *Information and Control*, vol. 8, no. 3, pp. 338–353, 1965. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S00199586590241X>
- [13] T. Obo and E. Yasuda, "Intelligent fuzzy controller for human-aware robot navigation," in *2018 12th France-Japan and 10th Europe-Asia Congress on Mechatronics*. IEEE, 2018, pp. 392–397.
- [14] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.
- [15] T. Fan, P. Long, W. Liu, and J. Pan, "Distributed multi-robot collision avoidance via deep reinforcement learning for navigation in complex scenarios," *The International Journal of Robotics Research*, vol. 39, no. 7, pp. 856–892, 2020.
- [16] S. Yao, G. Chen, Q. Qiu, J. Ma, X. Chen, and J. Ji, "Crowd-aware robot navigation for pedestrians with multiple collision avoidance strategies via map-based deep reinforcement learning," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 8144–8150.
- [17] Y. Liu, A. Xu, and Z. Chen, "Map-based deep imitation learning for obstacle avoidance," in *Proceedings of IROS-2018*, 2018, pp. 8644–8649.
- [18] L. Tai, J. Zhang, M. Liu, and W. Burgard, "Socially compliant navigation through raw depth inputs with generative adversarial imitation learning," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 1111–1117.
- [19] J. Ho and S. Ermon, "Generative adversarial imitation learning," *Advances in neural information processing systems*, vol. 29, pp. 4565–4573, 2016.
- [20] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," *Advances in neural information processing systems*, vol. 27, 2014.
- [21] T. Kruse, A. K. Pandey, R. Alami, and A. Kirsch, "Human-aware robot navigation: A survey," *Robotics and Autonomous Systems*, vol. 61, no. 12, pp. 1726–1743, 2013.
- [22] J. Borenstein and Y. Koren, "Real-time obstacle avoidance for fast mobile robots," *IEEE Transactions on systems, Man, and Cybernetics*, vol. 19, no. 5, pp. 1179–1187, 1989.
- [23] Borenstein, Johann and Koren, Yoram, "Real-time obstacle avoidance for fast mobile robots in cluttered environments," in *Proceedings., IEEE International Conference on Robotics and Automation*. IEEE, 1990, pp. 572–577.
- [24] J. Borenstein, Y. Koren *et al.*, "The vector field histogram-fast obstacle avoidance for mobile robots," *IEEE transactions on robotics and automation*, vol. 7, no. 3, pp. 278–288, 1991.
- [25] G. Ferrer, A. Garrell, and A. Sanfeliu, "Robot companion: A social-force based approach with human awareness-navigation in crowded environments," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2013, pp. 1688–1694.
- [26] G. Ferrer and A. Sanfeliu, "Behavior estimation for a complete framework for human motion prediction in crowded environments," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2014, pp. 5940–5945.
- [27] D. Mehta, G. Ferrer, and E. Olson, "Autonomous navigation in dynamic social environments using multi-policy decision making," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2016, pp. 1190–1197.
- [28] B. Kim and J. Pineau, "Socially adaptive path planning in human environments using inverse reinforcement learning," *International Journal of Social Robotics*, vol. 8, no. 1, pp. 51–66, 2016.
- [29] M. Everett, Y. F. Chen, and J. P. How, "Motion planning among dynamic, decision-making agents with deep reinforcement learning," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 3052–3059.
- [30] C. Chen, Y. Liu, S. Kreiss, and A. Alahi, "Crowd-robot interaction: Crowd-aware robot navigation with attention-based deep reinforcement learning," in *Proceedings of ICRA-2019*, 2019, pp. 6015–6022.
- [31] D. Dugas, J. Nieto, R. Siegart, and J. J. Chung, "Ian: Multi-behavior navigation planning for robots in real, crowded environments," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 11 368–11 375.
- [32] J. Z. Leibo, V. Zambaldi, M. Lanctot, J. Marecki, and T. Graepel, "Multi-agent reinforcement learning in sequential social dilemmas," *arXiv preprint arXiv:1702.03037*, 2017.
- [33] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [34] J. Redmon and A. Farhadi, "Yolov3: An incremental improvement," *arXiv preprint arXiv:1804.02767*, 2018.
- [35] T. Linder, S. Breuers, B. Leibe, and K. O. Arras, "On multi-modal people tracking from mobile platforms in very crowded and dynamic environments," in *Proceedings of ICRA-2016*. IEEE, 2016, pp. 5512–5519.