

L2B: Learning to Balance the Safety-Efficiency Trade-off in Interactive Crowd-aware Robot Navigation

Mai Nishimura¹ and Ryo Yonetani¹

Abstract—This work presents a deep reinforcement learning framework for interactive navigation in a crowded place. Our proposed *Learning to Balance (L2B)* framework enables mobile robot agents to steer safely towards their destinations by avoiding collisions with a crowd, while actively clearing a path by asking nearby pedestrians to make room, if necessary, to keep their travel efficient. We observe that the safety and efficiency requirements in crowd-aware navigation have a trade-off in the presence of social dilemmas between the agent and the crowd. On the one hand, intervening in pedestrian paths too much to achieve instant efficiency will result in collapsing a natural crowd flow and may eventually put everyone, including the self, at risk of collisions. On the other hand, keeping in silence to avoid every single collision will lead to the agent’s inefficient travel. With this observation, our L2B framework augments the reward function used in learning an interactive navigation policy to penalize frequent active path clearing and passive collision avoidance, which substantially improves the balance of the safety-efficiency trade-off. We evaluate our L2B framework in a challenging crowd simulation and demonstrate its superiority, in terms of both navigation success and collision rate, over a state-of-the-art navigation approach.

I. INTRODUCTION

We envision a future mobile robot system that can navigate crowded places, such as busy shopping malls and airports, as naturally as we do. Developing such an intelligent navigation system would enhance several practical applications including automated delivery services [1] and guidance at airports [2]. As shown in Fig. 1, to achieve this goal, we present a deep reinforcement learning (RL) framework for crowd-aware navigation, which enables agents to interact with a crowd not only by finding a bypass *safely* but also by actively clearing a path to arrive at their destinations *efficiently*.

Typically, it is easy for humans to navigate in congested environments safely and efficiently. For example, if someone cuts right in front of us, we simply stop walking to avoid potential collisions (*i.e.*, safe navigation). Also, if we are in a hurry to arrive at a destination in time, we typically call out to nearby people to make room (*i.e.*, efficient navigation). However, learning such advanced navigation skills from scratch is non-trivial because there exist no optimal solutions to determine when to avoid collisions passively or to clear a path actively. In the robotics domain, most prior works on mobile robot navigation have only focused on collision avoidance skills [3]–[6]. While they allow robots to move safely, too much collision avoidance also results in highly evasive maneuvers [5], [7]. On the other hand,




Fig. 1: Left: freezing robot problem, where a robot is struggling to find a bypass and as a result, takes unnecessary maneuvers. Right: our agent that can interact with a crowd by actively clearing a path or passively finding a bypass. Our goal is to enable such agents to balance the safety-efficiency trade-off in a highly congested environment.

actively addressing nearby people to move away, via sound notifications [2] or visual signs [8]–[10], would allow the robots to navigate efficiently by sticking with the original path plan. However, frequently making room in a crowd could collapse a natural crowd flow and may even increase the risk of collisions, especially in the extremely congested environments. Therefore, there is a trade-off between agent’s safety and efficiency in crowd-aware navigation, which begs our key question: “*how can agents learn to balance the safety-efficiency trade-off?*”

To address the above question, we develop a deep RL framework called *Learning to Balance (L2B)*. The proposed L2B framework enables crowd-aware navigation agents such as mobile robots to learn a hybrid policy allowing for choosing to either 1) seek a bypass (to passively avoid potential collisions with a crowd) or 2) actively address nearby persons, *e.g.*, by emitting a beeping sound [2] (to clear a planned path for a safe and efficient navigation). Our key insight is that the safety-efficiency trade-off may be viewed as a consequence of social dilemmas between a robot agent and a crowd. Specifically, they can both move safely and efficiently if they mutually cooperate and give way to each other, while doing otherwise will eventually result in navigation inefficiency (*i.e.*, sucker outcome) or unexpected collisions (*i.e.*, punishment from mutual defection). With this insight, we leverage the concept of Sequential Social Dilemmas (SSDs) [11] to augment the reward function used in learning the navigation policy, where the augmented reward function penalizes the undesirable outcomes and encourages mutual cooperation to balance travel safety and efficiency.

We evaluate the effectiveness of the L2B framework with

¹ Mai Nishimura and Ryo Yonetani are with OMRON SINIC X Corporation, Hongo 5-24-5, Bunkyo-ku, Tokyo, Japan. {mai.nishimura, ryo.yonetani}@sinicx.com

a simulation for challenging crowd-aware navigation tasks. Our experimental results demonstrated that the L2B outperformed the state-of-the-art SARL [12] navigation method. To the best of our knowledge, our work is the first to present a deep RL framework for learning a crowd-aware navigation policy that not only avoids collisions passively but also clears the path actively to make agent’s travels safe and efficient.

II. RELATED WORK

This section reviews related work mainly on 1) crowd-aware robot navigation, 2) modeling of social interactions in navigation tasks, and 3) modeling of crowd dynamics.

A. Crowd-aware Robot Navigation

A task of navigating robots in a crowd has been addressed based on reciprocal force models [13]–[15] or imitation learning [16]–[19]. Recently, deep RL approaches have achieved promising results in congested scenarios by jointly performing path planning and collision avoidance [6], [12], [20], [21]. All of these prior works, however, do not necessarily ensure the efficient travel of robots and can easily get trapped into a “freezing robot problem” [5], [7]. Specifically, once the environment surpasses a certain level of congestion, the planner decides that all forward paths are unsafe, and the robot freezes in place or takes unnecessary roundabout ways to avoid collisions. Despite much progress made to resolve this freezing problem [22], [23], it remains hard to enable robots to find a proper path in highly congested scenarios. Our work mitigates this problem by taking into account the robot’s ability to address surrounding pedestrians to actively clear a path.

B. Modeling Social Norms

Modeling and learning social norms, to allow robots to interact with humans properly, is also an active topic that has been studied in the context of navigation. Some work tried to learn “social etiquette” from pedestrian trajectories [24], which is then extended to robot motion planning [20]. Game-theoretic formulations are also studied widely [25]–[27] and extended to pedestrian modeling [28]. One important study that inspired our work is so-called Sequential Social Dilemmas (SSD) [11], [29], [30], which models how multiple agents cooperate in a complex pay-off structure with a Partially Observable Markov Decision Process (POMDP). Their multi-agent approach is confirmed effective in some Markov games with complex dilemmas such as the well-known Tragedy of the Commons [31]. In this work, we are interested in leveraging their formulation for interactive navigation tasks, where the key novelty is to design a reward function that takes into account the dilemma between a robot agent and a crowd to take a balance between navigation safety and efficiency.

C. Modeling Crowd Dynamics

Finally, modeling the dynamic behaviors of a crowd is crucial for both tasks of crowd-aware navigation and social interaction under crowded situations. Popular approaches include multi-agent interactions with reciprocal

force model [13], [14] and imitation learning [16]–[18]. These methods have been utilized not only in crowd simulations but also in controlling mobile robots navigating among people [32]. In our experiment, we will make use of the Emotional Reciprocal Velocity Obstacles (ERVO) model [33] for simulating a crowded environment, where each pedestrian will try to reach designated points or escape to safe places when they feel threatened by robot’s interventions.

III. BACKGROUND

A. Reinforcement Learning for Navigation

In this work, we consider a task of mobile robot navigation through a crowd, which we formulate as a sequential decision making problem with a POMDP. The robot and the crowd in the environment are regarded as two types of agents, namely **robot agent** and **crowd agent**, which are each driven by distinct policies $\pi, \tilde{\pi}$. Specifically, we regard multiple people in the environment as a single virtual agent to enable the modeling of a social dilemma between the robot and the crowd. This single crowd agent has a policy $\tilde{\pi}$ that is a fixed and unknown function. This allows us to formulate our problem as a standard crowd-aware navigation (same as [12]), where *only robot agent’s policy π is trainable, and the crowd agent is modeled as a part of the environment.*

At each time step t , the state \tilde{s}_t of the crowd agent is partially observable for the robot agent depending on its field of view and hidden goals of each constituent pedestrian in the crowd. Therefore, we describe the crowd agent’s state as a tuple $\tilde{s}_t = \langle \tilde{s}_t^o, \tilde{s}_t^h \rangle$, where \tilde{s}_t^o and \tilde{s}_t^h are observable (*e.g.*, locations of the pedestrian closest to the robot) and hidden parts of the \tilde{s}_t respectively. Accordingly, the robot agent observes a joint state $\mathbf{s}_t^{jn} = \langle s_t, \tilde{s}_t^o \rangle$ in each time step, where s_t is its own state such as the location of the self. Then the agent executes an action \mathbf{a}_t based on its policy π , and it will receive an instant reward $R(\mathbf{s}_t^{jn}, \mathbf{a}_t)$ designed to encourage the robot when reaching a goal and to penalize collisions with nearby pedestrians. The agent is then transitioned to the next state \mathbf{s}_{t+1}^{jn} based on the hidden state transition of crowd agent from $\langle \tilde{s}_t^o, \tilde{s}_t^h \rangle$ to $\langle \tilde{s}_{t+1}^o, \tilde{s}_{t+1}^h \rangle$.

With reinforcement learning, our objective is to obtain an optimal policy for the robot, $\pi^* : \mathbf{s}_t^{jn} \mapsto \mathbf{a}_t$ that maximizes the expectation of discounted total rewards. Following [12], one can also find the optimal value function V^* that encodes an estimate of the expected return as follows:

$$V^*(\mathbf{s}_t^{jn}) = \mathbb{E} \left[\sum_{t'=t}^T \gamma^{t'-v_{\text{pref}}} R(\mathbf{s}_{t'}^{jn}, \pi^*(\mathbf{s}_{t'}^{jn})) \right], \quad (1)$$

where $\gamma \in [0, 1)$ is a discount factor, and preferred velocity v_{pref} is used to normalize a term in the discount factor for numerical stability [6]. With the value iteration method, the optimal policy $\pi^*(\mathbf{s}_t^{jn})$ can then be derived as:

$$\pi^*(\mathbf{s}_t^{jn}) = \underset{\mathbf{a}_t}{\operatorname{argmax}} R(\mathbf{s}_t^{jn}, \mathbf{a}_t) + \gamma^{\Delta t \cdot v_{\text{pref}}} \int_{\mathbf{s}_{t+\Delta t}^{jn}} T(\mathbf{s}_t^{jn}, \mathbf{s}_{t+\Delta t}^{jn} | \mathbf{a}_t) V^*(\mathbf{s}_{t+\Delta t}^{jn}) d\mathbf{s}_{t+\Delta t}^{jn}. \quad (2)$$

To better deal with a high-dimensional state space, we approximate the value function V with a deep neural network.

B. Modeling Social Dilemmas

As we described earlier, we observe that a robot agent navigating in a crowd involves a type of social dilemmas. In the context of sequential decision making problems, Sequential Social Dilemmas (SSDs) deal with situations where individual agents are tempted to increase their own payoffs at the cost of lowering total rewards [11]. Specifically, SSDs consider a two-player partially observable Markov game with the joint policy denoted by $\vec{\pi} = (\pi, \tilde{\pi})$. The long-term payoff $V_i^{\vec{\pi}}(s_0)$ to player i from $t = 0$ can be represented as

$$V_i^{\vec{\pi}}(s_0) = \mathbb{E} \left[\sum_{t=0}^T \gamma^{t-v_{\text{pref}}} R(s_t, \vec{\pi}(s_t)) \right]. \quad (3)$$

Let π^c and π^d be cooperative and defecting policies in social dilemma games. In the SSDs, we consider four possible outcomes $\langle \mathcal{R}, \mathcal{P}, \mathcal{S}, \mathcal{T} \rangle := \langle V^{\pi^c, \pi^c}(s_0), V^{\pi^d, \pi^d}(s_0), V^{\pi^c, \pi^d}(s_0), V^{\pi^d, \pi^c}(s_0) \rangle$, where \mathcal{R} is a reward of mutual cooperation, \mathcal{P} is a punishment from mutual defection, \mathcal{S} is a sucker outcome for cooperation with a defecting partner, and \mathcal{T} is a temptation outcome when defecting against a cooperative partner. They satisfy the following inequalities; a) $\mathcal{R} > \mathcal{P}$ (mutual cooperation is better than mutual defection), b) $\mathcal{R} > \mathcal{S}$ (and is also better than being exploited), c) $\mathcal{T} > \mathcal{R}$ (exploiting the other is preferred to mutual cooperation), and d) $2\mathcal{R} > \mathcal{T} + \mathcal{S}$ (unilateral cooperation or defecting at equal probability is worse than mutual cooperation). In the next section, we augment our reward function by involving the SSD-like pay-off structure, which allows us to naturally balance the safety-efficiency trade-off in interactive navigation tasks.

IV. APPROACH

A. Social Dilemmas in Crowd-aware Navigation

For a robot agent that is capable of clearing a path actively, it is not necessarily optimal to do so whenever it finds someone on the path. Such actions can collapse a natural crowd flow, and may even lead to significant delays in reaching the goal or unexpected collisions if the place around the robot agent gets extremely crowded. On the other hand, keeping silence and avoiding every collision passively, which is what most of the existing approaches do, is inefficient especially in crowded scenes [23].

In the SSD's terminology, π^c can be viewed as a passive policy of one type of agent to avoid collisions by giving way to the other, whereas π^d is the contrary. Then, mutual cooperation $\mathcal{R} = V^{\pi=\pi^c, \tilde{\pi}=\pi^c}$ holds when robot and crowd agents give way to each other for making each travel reasonably safe and efficient. The robot agent suffers from navigation inefficiency due to frequent passive collision avoidance, which corresponds to $\mathcal{S} = V^{\pi=\pi^c, \tilde{\pi}=\pi^d}$. Although the robot agent may get close to the goal faster via active path clearing *i.e.*, $\mathcal{T} = V^{\pi=\pi^d, \tilde{\pi}=\pi^c}$, that will eventually result in mutual defection $\mathcal{P} = V^{\pi=\pi^d, \tilde{\pi}=\pi^d}$ once the crowd agent is no longer able to make room for the robot in a collapsed flow.

B. Learning to Balance the Safety-Efficiency Trade-off

With the above insight, we propose to better balance the navigation safety and efficiency by encouraging the robot to stay in \mathcal{R} while avoiding \mathcal{S} and \mathcal{T} that leads to \mathcal{P} . To incorporate this SSD structure into RL-based navigation tasks, we augment the reward function $R(s_t^{jn}, \mathbf{a}_t)$ to take into account the result of the robot's active path clearing:

$$R(s_t^{jn}, \mathbf{a}_t) = R_e(s_t^{jn}, \mathbf{a}_t) + R_s(s_t^{jn}, \mathbf{a}_t), \quad (4)$$

where R_e is the reward from environments, and R_s is that from the crowd agent. R_e is given by:

$$R_e(s_t^{jn}, \mathbf{a}_t) = \begin{cases} 1.0 - \alpha \frac{t}{t_{\text{lim}}} & \text{if } \mathbf{p}_t = \mathbf{p}^{(g)} \\ -0.25 & \text{elseif } d_t < d_{\text{min}} \\ 0 & \text{otherwise,} \end{cases} \quad (5)$$

where d_t is the distance between the robot agent and the crowd agent (*e.g.*, the closest pedestrian) at time t and t_{lim} is a time limit for the task completion. \mathbf{p}_t and $\mathbf{p}^{(g)}$ are the position of the robot at time t and its destination, respectively. This reward will monotonically decay over time so that the robot agent will be encouraged to reach the goal as early as possible.

On the other hand, We define R_s as follows:

$$R_s(s_t^{jn}, \mathbf{a}_t) = \begin{cases} \beta(d_t - r^{(b)}) & \text{if } d_t < r^{(b)} \wedge b_t = 1 \\ \eta(d_t - d_{\text{disc}}) & \text{elseif } d_t < d_{\text{disc}} \\ 0 & \text{otherwise,} \end{cases} \quad (6)$$

where $r^{(b)}$ is the effective range of active path clearing actions within which persons in the crowd will get influenced and move away, and d_{disc} is the minimum discomfort distance between the robot agent and the crowd agent set to encourage earlier collision avoidance. A binary vector $b_t \in \{0, 1\}$ is 1 when path clearing action is invoked by a robot and 0 otherwise. β and η are hyper-parameters that respectively adjust the influence of active path clearing and a penalty due to crowd agent's discomfort. The first term with β represents the robot's aggressiveness in path clearing strategy, where a robot agent is tempted to clear a path rather than finding a bypass, *i.e.*, the transition from \mathcal{R} to \mathcal{T} . However, too much active path clearing will harm the crowd flow, *i.e.*, the transition from \mathcal{T} to \mathcal{P} . On the other hand, the second term with η corresponds to \mathcal{R} to \mathcal{S} ; permitting crowd agent's free movements via collision avoidance too much will make the robot travel inefficient. Moreover, we set β and η to satisfy $\eta > \beta$; the robot agent will be penalized more heavily if it gets too close to nearby persons than actively clearing a path, *i.e.*, $\mathcal{T} > \mathcal{S}$. A high return by \mathcal{R} is obtained when the robot agent reaches its destination as early as possible to increase R_e while avoiding $\mathcal{T}, \mathcal{S}, \mathcal{P}$ by keeping the effect from R_s as small as possible. Doing so corresponds to reasonably reducing the frequency of both crowd avoidance and active path clearing; *i.e.*, balancing the safety-efficiency trade-off.

With this reward function, the value function can be trained via a standard V-learning based on a temporal

difference method with experience replay and fixed target techniques [34], such as done in [12].

V. SIMULATION

To evaluate the effectiveness of the proposed L2B framework, we develop a simulation environment on top of the OpenAI Gym [35] for challenging interactive crowd-aware navigation tasks.

A. Environment Setup

Following [6], [12], we implement circle crossing scenarios where N pedestrians and the robot are positioned randomly on a circle with a certain fixed radius, which will each walk to their destination set at the opposite side of the circle (see also Fig. 3). Random noises are added to the initial locations and destinations to make scenes further diverse.

For the i -th pedestrian, let $\mathbf{p}_t^{(i)} \in \mathbb{R}^2$ and $\mathbf{v}_t^{(i)} \in \mathbb{R}^2$ be the location and velocity at time t . Similarly, let $\mathbf{p}_t \in \mathbb{R}^2$ and $\mathbf{v}_t \in \mathbb{R}^2$ be those of the robot agent. We also denote the sizes of the pedestrian and the robot by $r^{(i)}, r^{(c)} \in \mathbb{R}_+$, and the distance between the i -th pedestrian and the robot at t by $d_t^{(i)} = \|\mathbf{p}_t^{(i)} - \mathbf{p}_t\| \in \mathbb{R}_+$. They are judged to collide with each other if $r^{(c)} + r^{(i)} \geq d_t^{(i)}$. As for the robot's ability to actively clear its own path, let $b_t \in \{0, 1\}$ be a binary vector indicating if an active path clearing action (which we also refer to as a *beep* action in our experiment section by following [2]) is executed at time t , and $r^{(b)} \in \mathbb{R}_+$ be its effective range. Finally, let $\mathbf{p}^{(g)}$ be the robot's destination, $d_t^{(g)} = \|\mathbf{p}_t - \mathbf{p}^{(g)}\|$ be the direct distance from the current robot's location to the destination, and $v_{\text{pref}} \in \mathbb{R}_+$ be a preferred travel speed of the robot.

With the above notations, a state of the robot agent \mathbf{s}_t and that of the crowd agent $\tilde{\mathbf{s}}_t^o$ are defined as follows:

$$\mathbf{s}_t = [d_t^{(g)}, \mathbf{v}_t, v_{\text{pref}}, r^{(c)}, r^{(b)}], \quad (7)$$

$$\tilde{\mathbf{s}}_t^o = [\tilde{\mathbf{p}}_t, \tilde{\mathbf{v}}_t, \tilde{d}_t, \tilde{r}], \quad (8)$$

where $\tilde{\mathbf{p}}_t = \mathbf{p}_t^{(j)}$, $j = \arg \min_i d_t^{(i)}$, is the location of pedestrian located closest to the robot at time t , and the same applies to $\tilde{\mathbf{v}}_t, \tilde{d}_t, \tilde{r}$. Also, the robot agent's action \mathbf{a}_t is defined using $\mathbf{v}_{t+1}, b_{t+1}$ (will be given concretely in the next section), assuming that the velocity of the robot can be controlled instantly.

B. Modeling Reactions of Pedestrians

One critical choice of designs for simulating interactive navigation scenarios is how a crowd reacts to active path clearing by the robot. In this work, we implement a simplified version of Emotional Reciprocal Velocity Obstacles (ERVO) [33], which is an extension of RVO considering a person's emotional reaction towards a threat. ERVO simulates how people in a panic choose their paths to safe places or a planned goal in a realistic way.

Formally, let $\Gamma(\mathbf{p}_t^{(i)})$ be the degree of influence of active path clearing for a pedestrian at location $\mathbf{p}_t^{(i)}$. We represent




Fig. 2: ERVO-based reactive pedestrians. A circle and triangles represent a robot with size $r^{(c)}$ and pedestrians with size $r^{(i)}$, respectively (Left). When the robot executes an active path clearing action, all the pedestrians within the half-circle with radius $r^{(b)}$ (Right) are affected to change their directions based on Eq. (10).

TABLE I: Quantitative results. Success rate, collision rate, timeout rate, and average navigation time for the proposed L2B-SARL and the baseline SARL under different numbers of pedestrian N in a crowd.

| Methods | N | Success | Collision | Timeout | Time |
|-----------------|---------|--------------|--------------|--------------|--------------|
| L2B-SARL | 20 | 0.906 | 0.094 | 0.000 | 13.85 |
| SARL [12] | | 0.700 | 0.184 | 0.116 | 13.50 |
| L2B-SARL | 15 | 0.880 | 0.118 | 0.002 | 11.60 |
| SARL [12] | | 0.778 | 0.064 | 0.158 | 12.43 |
| L2B-SARL | 10 | 0.904 | 0.086 | 0.004 | 11.31 |
| SARL [12] | | 0.922 | 0.046 | 0.032 | 11.91 |
| L2B-SARL | 5 | 0.978 | 0.020 | 0.002 | 10.14 |
| SARL [12] | | 0.966 | 0.032 | 0.002 | 10.09 |
| L2B-SARL | average | 0.917 | 0.079 | 0.002 | 11.72 |
| SARL [12] | | 0.841 | 0.081 | 0.077 | 11.98 |

this influence with a Gaussian distribution such as

$$\Gamma(\mathbf{p}_t^{(i)}) = \begin{cases} \frac{1}{\sqrt{2\pi}r^{(b)}} \exp\left(-\frac{(\mathbf{p}_t^{(i)} - \mathbf{p}_t)^2}{2r^{(b)2}}\right) & \text{if } d_t^{(i)} < r^{(b)} \wedge b_t = 1 \\ 0 & \text{otherwise.} \end{cases} \quad (9)$$

Then, the i -th agent will change its velocity $\mathbf{v}_t^{(i)}$ as follows:

$$\mathbf{v}_t^{(i)} \leftarrow \begin{cases} \Gamma(\mathbf{p}_t^{(i)}) \cdot \frac{\mathbf{p}_t^{(i)} - \mathbf{p}_t}{d_t^{(i)}} & \text{if } d_t^{(i)} < r^{(b)} \wedge b_t = 1 \\ \mathbf{v}_t^{(i)} & \text{otherwise.} \end{cases} \quad (10)$$

As illustrated in Fig. 2, the agents within r_b will act to escape from the path clearing influence that attenuates based on the Gaussian distribution. On the basis of these fundamental principles of ERVO, we extend a collision avoidance simulation called Optimal Reciprocal Collision Avoidance (ORCA) [14] to synthesize realistic flows of a crowd.

VI. EVALUATIONS

With the simulation introduced above, we evaluate a performance of a state-of-the-art crowd-aware navigation method called SARL [12], and its extended version equipped with the active path clearing ability trained in the proposed L2B framework, which we will refer to as **L2B-SARL**.

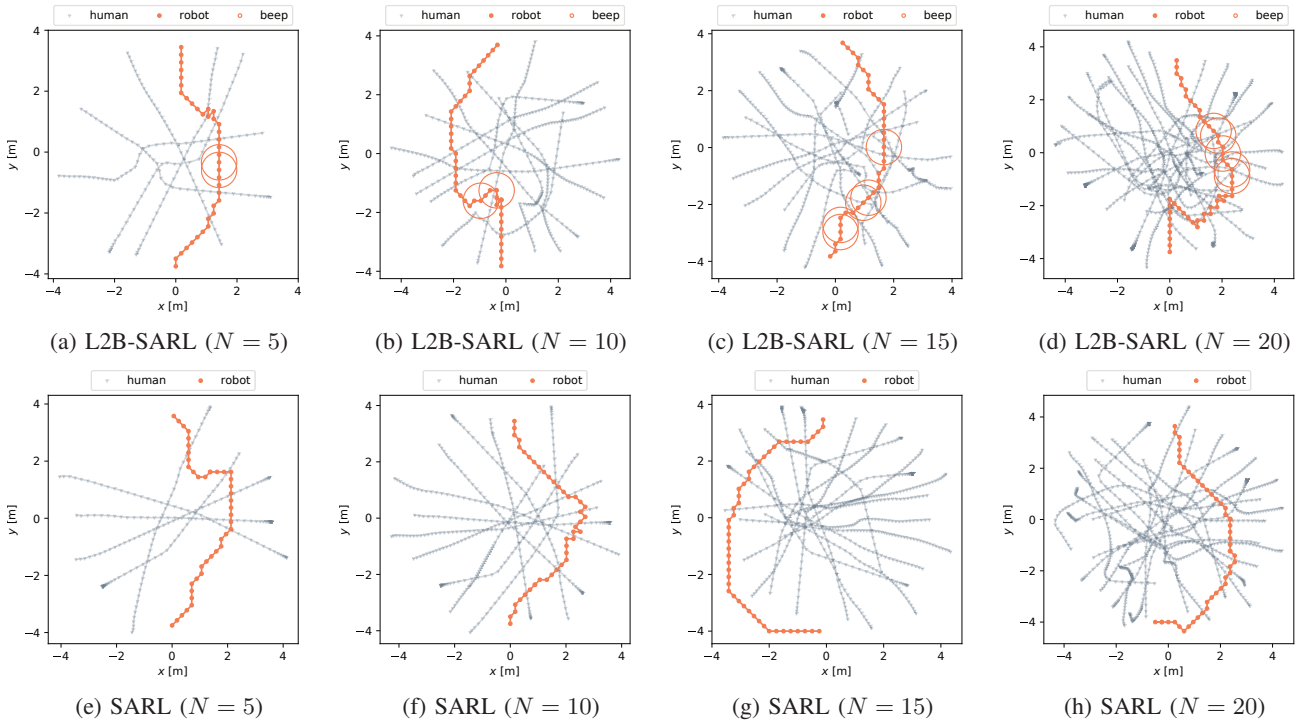


Fig. 3: Qualitative results. (a)-(d) show results of our proposed method L2B-SARL, and (e)-(h) show those of SARL. Gray points represent human agents, and filled orange circle represents robot agents. Bigger orange circle represent path clearing actions.

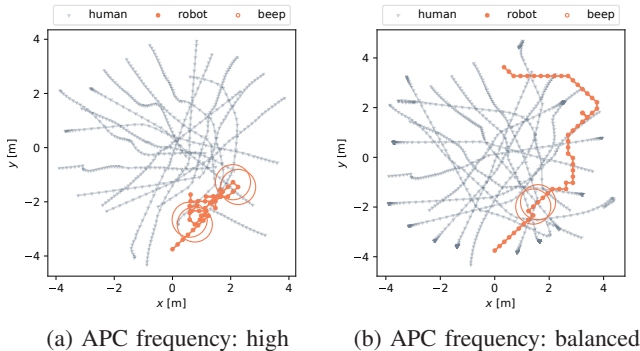


Fig. 4: Effect of path clearing frequency. (a) When β is set to be small, too many active path clearing (APC) caused traffic confusions and increased the possibility of collisions. (b) Our approach can learn to choose path clearing and collision avoidance adequately.

A. Experimental Setup

1) *Environments*: With our simulation, we synthesize diverse crowd-aware navigation tasks with the number of pedestrians $N \in \{5, 10, 15, 20\}$. Kinematics of the robot agent is assumed to be holonomic, *i.e.*, it can move in any direction without spin. The velocity of robot, v_t , was discretized in two speeds in $\{0, v_{\text{pref}}\}$ and eight orientations spaced evenly between $[0, 2\pi)$. In total, there were 9-dimensional action spaces for SARL (*i.e.*, moving one of the eight directions with v_{pref} or standing still), and 17-

dimensional action spaces (eight directions \times with/without path-clearing beeps, and standing still) for the L2B-SARL. The preferred speed of the agent v_{pref} was set to 1.0m/s. For the proposed reward function in Section IV-B, we set discomfort distance d_{disc} to be 0.2m, the effective range of path clearing actions $r^{(b)}$ to be 1.0m, and the diameter of both the robot and the human agents $r^{(c)}$, $r^{(i)}$ to be 0.3m. The attenuating reward coefficient for reaching a goal α and discomfort penalty factor η are configured to be 0.1 and 0.5, respectively. For β , we set $\beta = 0.2$ unless specified otherwise. In the simulation, the radius of circle on which pedestrians and the robot were placed, was set to 4.0m.

2) *Training Details*: For both L2B-SARL and SARL, we implemented a value function with the attention-based network proposed in [12], which allows agents to observe nearby pedestrians effectively. Intuitively, the network models pairwise human-robot interactions explicitly while encoding human-human interactions in a coarse-grained feature map, and aggregates the interactions by a self-attention mechanism. This value network was first initialized via imitation learning from 3k episodes collected from the ORCA [14] policy, using the Adam optimizer [36] with the learning rate of 0.01 for 50 epochs. Then, we train the value network in an RL loop via the Adam optimizer with learning rate 0.001, mini-batch of size 100, and discount factor $\gamma = 0.9$, for

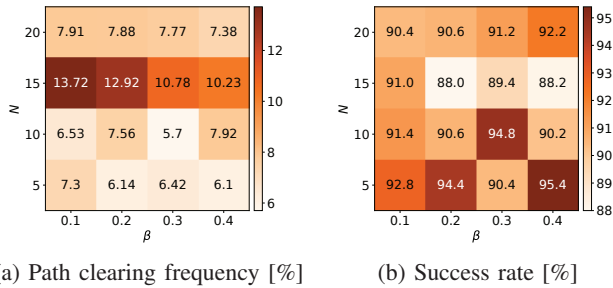


Fig. 5: Effect of β . (a) Frequency of active path-clearing actions and (b) success rate of episodes for different combination of penalty factor β and the number of pedestrians in a crowd N .

20k episodes¹. We adopted a standard ϵ -greedy exploration scheme where ϵ decayed linearly from 0.5 to 0.1 in the first 5k episodes and stayed at 0.1 for the rest.

3) *Evaluation Scheme*: For each environment with the number of people in crowd $N \in \{5, 10, 15, 20\}$, we evaluated the trained models with 500 random test cases with the following metrics; “Success”: the rate of robot reaching its goal without a collision, “Collision”: the rate of robot terminating its navigation due to collisions with other humans, “Timeout”: the rate of robot unable to reach a goal within time limit t_{lim} , and “Time”: the average robots navigation time to reach its goal in seconds. For each evaluation, the random seed was set to be the same so that each test case can be evaluated on exactly the same sequence.

B. Results

1) *Quantitative Evaluation*: Table I summarizes quantitative evaluation results. Overall, we confirmed that L2B-SARL demonstrated high success rates regardless of the number of people N , whereas the baseline SARL degraded its performance as the environment got more crowded. This is mainly because the baseline agent could only find a bypass passively when it found someone on its path, resulting in high timeout rates. On the other hand, the L2B-SARL provides an extremely low timeout rate thanks to its ability to actively clear a path, at the small cost of small increases of collision rates. The average navigation time for the successful sequences was comparable between L2B-SARL and SARL, These results show that the proposed L2B framework allows us to better balance the safety and efficiency trade-off in crowd-aware navigation tasks.

2) *Effect of Active Path Clearing*: We further investigated how the robot learned to actively clear a path adequately under several different choices of β . Fig. 5 summarizes frequencies of path clearing actions (a) and success rates (b) for $\beta \in \{0.1, 0.2, 0.3, 0.4\}$. Importantly, there was no monotonic tendencies for the path clearing frequency for

¹Note that SARL performed poorly in environments with $N = 20$, which reached only 0.04 success rate. Therefore firstly we trained the policy in environments with $N = 10$ for the first 10k episodes and then transferred to $N = 20$ environments for the remaining 10k episodes.

any choice of β . The robot is learned to actively clear a path more frequently as N becomes larger by 5 – 15, but in highly congested scenarios at $N = 20$, kept its frequency lower than $N = 15$. With constantly high success rates shown in (b), these results indicate that the robot could be aware that the ideal path clearing frequency maximizes the cumulative reward would depends on a congestion factor. For less congested environments at $N = 5, 10$, the frequency is preferred to be low because there are enough spaces to find a bypass as shown in Fig. 3(a). For more congested environments at $N = 15$, the success rate increases as β becomes lower because frequent path clearing allows agents to reach a goal with a shorter path than when only performing collision avoidance as shown in Fig. 3(c)(g). However, at highly congested environments $N = 20$, the frequency is preferred to be low again and the success rate decreases as β becomes lower. One possible interpretation of this result is, in highly congested scenarios as in Fig. 3(d), too much active path clearing will also make its travel inefficient or unsafe in the presence of social dilemmas against the crowd.

C. Qualitative Evaluation

Fig. 3 visualizes some typical results for L2B-SARL in (a)-(d) and SARL in (e)-(h). We observe that SARL’s paths became roundabout and inefficient by trying to bypass crowded regions at the center of the environment in $N = 15, 20$. In contrast, the L2B-SARL agent was able to navigate through a crowd by actively clearing its path (denoted by circles). Moreover, Fig. 4 (a) shows the special case that the agent is tempted to clear a path frequently and the episode ended with collision. It indicates that too much path clearing does harm a robot’s efficient travel, as doing so is not effective when pedestrians being addressed had little room to move away. If the path clearing frequency is balanced by our learnt adaptive policy, the agent can efficiently reach a goal with the moderate number of times of path clearing as shown in Fig. 4(b).

VII. CONCLUSION

We have presented a new deep RL framework called L2B for crowd-aware navigation tasks, which enabled robotic agents to navigate through a crowd safely and efficiently. The key idea is to equip the agents with the ability to actively clear a path as well as passively find a bypath to avoid collisions. With a reward function that takes into account the presence of social dilemmas between the robot and a crowd, the proposed L2B framework allows us to learn a navigation policy to choose these two actions adequately to take a good balance between travel safety and efficiency. Our extensive simulation experiments demonstrated the superiority of the proposed approach over a state-of-the-art navigation method.

Currently, we limit our study to assume that all the pedestrians only react passively based on the fixed policy. One interesting extension of the proposed work is to involve pedestrians who also try to clear a path actively. That will make more explicit the presence of a social dilemma

structure in a crowd flow, also requiring a new technique for simulating such active pedestrians in a realistic scenario. Another possible direction for future work is to formulate this crowd-aware navigation task in a multi-agent RL problem, where each pedestrian in a crowd also allowed to improve its policy to better cooperate with the robot. Such a direction is beneficial for practical robotics applications such as swarm robotics [37] and multiple vehicle control [38].

ACKNOWLEDGMENTS

The authors would like to thank Mohammadamin Barekatin, Tatsunori Tanai and Yoshihisa Ijiri for the insightful discussions and helpful feedback on the manuscript.

REFERENCES

- [1] D. Feil-Seifer and M. J. Mataric, "Socially assistive robotics," *IEEE Robotics & Automation Magazine*, vol. 18, no. 1, pp. 24–31, 2011.
- [2] S. Kayukawa, K. Higuchi, J. Guerreiro, S. Morishima, Y. Sato, K. Kitani, and C. Asakawa, "Bbeep: A sonic collision avoidance system for blind travellers and nearby pedestrians," in *CHI Conference on Human Factors in Computing Systems*, 2019, pp. 1–12.
- [3] J. Borenstein and Y. Koren, "Real-time obstacle avoidance for fast mobile robots in cluttered environments," in *International Conference on Robotics and Automation*, 1990, pp. 572–577.
- [4] D. Fox, W. Burgard, and S. Thrun, "The dynamic window approach to collision avoidance," *IEEE Robotics & Automation Magazine*, vol. 4, no. 1, pp. 23–33, 1997.
- [5] P. Trautman and A. Krause, "Unfreezing the robot: Navigation in dense, interacting crowds," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2010, pp. 797–803.
- [6] Y. F. Chen, M. Liu, M. Everett, and J. P. How, "Decentralized non-communicating multiagent collision avoidance with deep reinforcement learning," in *IEEE International Conference on Robotics and Automation*, 2017, pp. 285–292.
- [7] P. Trautman, J. Ma, R. M. Murray, and A. Krause, "Robot navigation in dense human crowds: Statistical models and experimental studies of human-robot cooperation," *International Journal of Robotics Research*, vol. 34, no. 3, pp. 335–356, 2015.
- [8] T. Matsumaru, T. Kusada, and K. Iwase, "Mobile robot with preliminary-announcement function of forthcoming motion using light-ray," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2006, pp. 1516–1523.
- [9] T. Matsumaru, "Mobile robot with preliminary-announcement and indication function of forthcoming operation using flat-panel display," in *IEEE International Conference on Robotics and Automation*, 2007, pp. 1774–1781.
- [10] A. Watanabe, T. Ikeda, Y. Morales, K. Shinozawa, T. Miyashita, and N. Hagita, "Communicating robotic navigational intentions," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2015, pp. 5763–5769.
- [11] J. Z. Leibo, V. Zambaldi, M. Lanctot, J. Marecki, and T. Graepel, "Multi-agent reinforcement learning in sequential social dilemmas," in *International Conference on Autonomous Agents and Multi-Agent Systems*, 2017, pp. 464–473.
- [12] C. Chen, Y. Liu, S. Kreiss, and A. Alahi, "Crowd-robot interaction: Crowd-aware robot navigation with attention-based deep reinforcement learning," in *IEEE International Conference on Robotics and Automation*, 2019, pp. 6015–6022.
- [13] J. Van den Berg, M. Lin, and D. Manocha, "Reciprocal velocity obstacles for real-time multi-agent navigation," in *IEEE International Conference on Robotics and Automation*, 2008, pp. 1928–1935.
- [14] 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.
- [15] G. Ferrer, A. Garrell, and A. Sanfeliu, "Robot companion: A social-force based approach with human awareness-navigation in crowded environments," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2013, pp. 1688–1694.
- [16] L. Tai, J. Zhang, M. Liu, and W. Burgard, "Socially compliant navigation through raw depth inputs with generative adversarial imitation learning," in *IEEE International Conference on Robotics and Automation*, 2018, pp. 1111–1117.
- [17] 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.
- [18] Y. Liu, A. Xu, and Z. Chen, "Map-based deep imitation learning for obstacle avoidance," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2018, pp. 8644–8649.
- [19] M. Pfeiffer, U. Schwesinger, H. Sommer, E. Galceran, and R. Siegwart, "Predicting actions to act predictably: Cooperative partial motion planning with maximum entropy models," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2016, pp. 2096–2101.
- [20] Y. F. Chen, M. Everett, M. Liu, and J. P. How, "Socially aware motion planning with deep reinforcement learning," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2017, pp. 1343–1350.
- [21] M. Everett, Y. F. Chen, and J. P. How, "Motion planning among dynamic, decision-making agents with deep reinforcement learning," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2018, pp. 3052–3059.
- [22] T. Fan, X. Cheng, J. Pan, P. Long, W. Liu, R. Yang, and D. Manocha, "Getting robots unfrozen and unlost in dense pedestrian crowds," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1178–1185, 2019.
- [23] Y. Chen, C. Liu, M. Liu, and B. E. Shi, "Robot navigation in crowds by graph convolutional networks with attention learned from human gaze," *arXiv preprint arXiv:1909.10400*, 2019.
- [24] A. Robicquet, A. Sadeghian, A. Alahi, and S. Savarese, "Learning social etiquette: Human trajectory understanding in crowded scenes," in *European Conference on Computer Vision*, 2016, pp. 549–565.
- [25] C. Claus and C. Boutilier, "The dynamics of reinforcement learning in cooperative multiagent systems," in *AAAI Conference on Artificial Intelligence*, vol. 1998, no. 746-752, 1998, p. 2.
- [26] M. Tan, "Multi-agent reinforcement learning: Independent vs. cooperative agents," in *International Conference on Machine Learning*, 1993, pp. 330–337.
- [27] J. F. Fisac, E. Bronstein, E. Steffansson, D. Sadigh, S. S. Sastry, and A. D. Dragan, "Hierarchical game-theoretic planning for autonomous vehicles," in *IEEE International Conference on Robotics and Automation*, 2019, pp. 9590–9596.
- [28] W.-C. Ma, D.-A. Huang, N. Lee, and K. M. Kitani, "Forecasting interactive dynamics of pedestrians with fictitious play," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 774–782.
- [29] N. Jaques, A. Lazaridou, E. Hughes, C. Gulcehre, P. Ortega, D. Strouse, J. Z. Leibo, and N. De Freitas, "Social influence as intrinsic motivation for multi-agent deep reinforcement learning," in *International Conference on Machine Learning*, 2019, pp. 3040–3049.
- [30] M. Lanctot, V. Zambaldi, A. Gruslys, A. Lazaridou, K. Tuyls, J. Pérolat, D. Silver, and T. Graepel, "A unified game-theoretic approach to multiagent reinforcement learning," in *Advances in Neural Information Processing Systems*, 2017, pp. 4190–4203.
- [31] G. Hardin, "The tragedy of the commons," *Science*, vol. 162, no. 3859, pp. 1243–1248, 1968.
- [32] T. Fraichard and V. Levesy, "From crowd simulation to robot navigation in crowds," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 729–735, 2020.
- [33] 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.
- [34] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski et al., "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [35] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, "Openai gym," 2016.
- [36] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [37] M. Brambilla, E. Ferrante, M. Birattari, and M. Dorigo, "Swarm robotics: a review from the swarm engineering perspective," *Swarm Intelligence*, vol. 7, no. 1, pp. 1–41, 2013.
- [38] Y. Cao, D. Stuart, W. Ren, and Z. Meng, "Distributed containment control for multiple autonomous vehicles with double-integrator dynamics: algorithms and experiments," *IEEE Transactions on Control Systems Technology*, vol. 19, no. 4, pp. 929–938, 2010.