

# SocNavGym: A Reinforcement Learning Gym for Social Navigation

Aditya Kapoor<sup>\*1</sup>, Sushant Swamy<sup>\*2</sup>, Pilar Bachiller<sup>3</sup> and Luis J. Manso<sup>4</sup>

**Abstract**—It is essential for autonomous robots to be socially compliant while navigating in human-populated environments. Machine Learning and, especially, Deep Reinforcement Learning have recently gained considerable traction in the field of Social Navigation. This can be partially attributed to the resulting policies not being bound by human limitations in terms of code complexity or the number of variables that are handled. Unfortunately, the lack of safety guarantees and the large data requirements by DRL algorithms make learning in the real world unfeasible. To bridge this gap, simulation environments are frequently used. We propose SocNavGym, an advanced simulation environment for social navigation that can generate a wide variety of social navigation scenarios and facilitates the development of intelligent social agents. SocNavGym is lightweight, fast, easy to use, and can be effortlessly configured to generate different types of social navigation scenarios. It can also be configured to work with different hand-crafted and data-driven social reward signals and to yield a variety of evaluation metrics to benchmark agents' performance. Further, we also provide a case study where a Dueling-DQN agent is trained to learn social-navigation policies using SocNavGym. The results provide evidence that SocNavGym can be used to train an agent from scratch to navigate in simple as well as complex social scenarios. Our experiments also show that the agents trained using the data-driven reward function display more advanced social compliance in comparison to the heuristic-based reward function.

## I. INTRODUCTION

Social compliance is key to deploying robots, not only for pedestrians' comfort but also to enhance efficiency. The sheer amount of research in the field is evidence of its importance [1]. Until recently, the vast majority of Social Navigation (SN) algorithms were hand-crafted and coded by humans. One of the limitations of this approach is its poor scalability in terms of the number of variables considered, which is partially because the equations involved become complex to understand, implement and perhaps most importantly, debug. Other factors that make incorporating other variables challenging include the lack of a principled theory of pedestrians' comfort, and the fact that the state of the world is only partially observable (e.g., pedestrians' hidden goals, intents, mood). Because of these limitations, the community is experiencing a strong trend toward data-driven methods that can account for more variables. The most

common SN approach making use of learning involves using hand-crafted algorithms in conjunction with learning-based human trajectory predictors. Examples of these works are [2], [3], [4], [5], [6], [7], where actions are planned using search trees and other heuristic-based algorithms relying on estimated human states [8]. Other learning-based works include social compliance estimators, such as the Social Navigation Graph Neural Network (SNGNN) [9]. A third family of promising social navigation algorithms that has gained recent popularity, are those using Deep Reinforcement Learning (DRL) and Learning from Demonstrations (LfD) [10], [11], [12], [13], [14], [15].

DRL systems have three especially critical components: 1) a DRL algorithm, 2) a neural architecture, and 3) an environment providing observations and rewards. While research in DRL fundamentally focuses on the first two aspects, a limitation in the current literature is that none of the reward functions used in the existing gym environments has been designed to reflect overall users' opinions about the contextual social compliance of the agent. The ones in use are more suitable to solve multi-agent collision avoidance rather than social navigation, which is a different, relatively easier problem. To the best of our knowledge, the reward functions used in the DRL-based SN literature are all piecewise defined functions where the piece considering the social compliance of the agent is a simple linear equation or a similar profile function that only consider pedestrians' positions. We strongly argue that, to learn socially-compliant policies via DRL, the reward function should be carefully designed to consider users' opinions.

Inverse Reinforcement Learning (IRL) has shown promising results in social navigation [16], [17], [18]. If trained with adequate data, a learned reward function can help tackle these issues, as a DRL architecture could be trained to optimise for it, enabling socially-compliant policy learning. The resulting behaviour could theoretically be as principled and effective as such a learned reward. Our hypothesis in this vein is that using a learned model for comfort as a reward function will lead to better social compliance than other simplistic or hand-crafted rewards. Therefore, our gym environment -SocNavGym- integrates SNGNN-v2 [19] in its reward function, the only model that we have found in the literature providing step-wise social compliance scores.

The desiderata established for SocNavGym's design is as follows:

- 1) be lightweight, fast and easy to use and configure, with as few software dependencies as possible;
- 2) simulate a rich variety of realistic social scenarios so that agents can be used in as many real-life settings as

<sup>1</sup>Aditya Kapoor is with Tata Consultancy Services, Research & Innovation, India. [aditya.kapoor1@tcs.com](mailto:aditya.kapoor1@tcs.com)

<sup>2</sup>Sushant Swamy is with Birla Institute of Technology and Science, Goa, India. [f20191031@goa.bits-pilani.ac.in](mailto:f20191031@goa.bits-pilani.ac.in)

<sup>3</sup>Pilar Bachiller is with Computer and Communication Technology Department, Universidad de Extremadura, Spain. [pilarb@unex.es](mailto:pilarb@unex.es)

<sup>4</sup>Luis J. Manso is with Autonomous Robotics and Perception Laboratory, Computer Science Department, Aston University, UK. [l.manso@aston.ac.uk](mailto:l.manso@aston.ac.uk)

\* is for equal contribution.

possible;

- 3) be easy to modify and extend;
- 4) provide a high-quality social reward to ensure that near-optimal policies meet real pedestrians' expectations; and
- 5) provide established evaluation metrics to benchmark agents' performance.

To the best of our knowledge, no existing navigation gym satisfies all the aforementioned requirements (see Section II for details).

SocNavGym is a 2D simulation environment built by extending the OpenAI's MultiAgent Particle Environment (MPE) [20]. We take inspiration from the existing work in the literature and incorporate some of the most relevant features into SocNavGym. Examples of additional features include static obstacle-like entities, static and dynamic agent formations, and interactions between different entities (*e.g.*, human-looking-at-laptop, human-talking-with-human). We also integrate SNGNN-v2 [19], a data-driven reward function trained on a users' opinions dataset [21], [22]. Further, we train a DuelingDQN [23] model to show that our simulation environment can be used to train DRL algorithms in a variety of rich social situations.

## II. RELATED WORK

Many high-fidelity simulators have been developed to collect data and train models for robotic applications *e.g.*, CARLA [24], iGibson [25]. Crowd simulation platforms like Nomad [26] and Menge [27] are designed to model individuals and crowds. On the contrary, SEAN [28] and SocialGym [29] are designed to train and evaluate social agents but their reliance on Unity [30] and Robot Operating System (ROS), might not be ideal for many DRL researchers due to additional complexity and computational resources required. Although these provide good physics simulations, they require a considerable amount of computing power that does not seem justifiable for this specific application. Also, to focus on decision-making for social navigation, one can abstract away the perception task that further weakens the need for high-fidelity simulators. SocNavGym and other state-of-the-art SN gyms (*e.g.*, SARL [10], DSRNN [11]), use low-fidelity simulators but, unlike SocNavGym, others fail to recreate realistic complex social settings. In SARL [10], scenarios are composed of humans forming a circle moving in the diametrically opposite direction using ORCA policy in an otherwise empty space. In DSRNN [11], static and dynamic groups are also present, but they still greatly differ from typical human-populated environments. SocialGym [29] goes a step further to introduce room-like scenarios with corridors and more realistic pedestrian movement. However, its dependence on ROS [31] creates additional overhead, and its setup and installation process is not as easy as desired. Further, bringing about any change in the environment requires the user to know the ROS framework which is not ideal. Additionally, none of these environments provides a principled social reward function that factors in human opinion, which makes agents trained

with them, treat humans as dynamic obstacles rather than social beings.

Another gap found in simulation environments is the lack of support for common interactions, such as those among humans and with other inanimate entities. Such interactions are important while training agents for realistic scenarios since the level of disturbance that an agent can cause depends on these interactions. For example, an agent could cause considerable disturbance to two people talking (human-human interaction) even if the agent does not enter the personal space of the individuals involved. These interactions should be taken into consideration, regardless of them being directly observable by the agents or expected to be inferred (SocNavGym supports both options).

With the exception of SocialGym [29] and SEAN [28], none of the reviewed SN gyms provide sufficient evaluation metrics (*e.g.*, time taken to reach the goal, minimum distance to human, discomfort level) to benchmark social performance. Even though these gyms provide a subset of the evaluation metrics, they do not provide any metric or reward function accounting for efficiency and human comfort.

## III. GYM DESIGN DECISIONS AND FEATURES

SocNavGym follows the standard OpenAI Gym API convention. Installation and setup are also straightforward, either from source or via pip. All adjustable parameters are described in the project's documentation and can be modified via configuration files without accessing the codebase. The project includes several sample configuration files to facilitate the use of the proposed tool. Moreover, it provides a DuelingDQN [23] baseline for reference, which is described in Sec. IV. We put forward a new open-source<sup>1</sup> gym that meets the desiderata described in Sec. II. In particular, SocNavGym provides the following features:

- 1) lightweight randomised scenario generation and simulation,
- 2) support for static objects, corridors and walls to bound rooms of different shapes and sizes (Fig. 1.i-j)),
- 3) creation and dispersion of group formations and interactions (*e.g.*, human-looking-at-laptop, group of humans) (Fig. 1 (a & e), (b & f), (c & g), (d & h)),
- 4) support for holonomic and non-holonomic robots that can be controlled with discrete or continuous action spaces,
- 5) integration of a learned reward function combining social and functional aspects (SNGNN-v2 [19]), simple heuristic-based reward functions used in other works [11], and an API to create custom reward functions,
- 6) it provides a variety of additional evaluation metrics to benchmark agents' performance (in particular the metrics suggested in the guidelines paper [32]),
- 7) it includes gaze as a human property that can be added to the observation space and allows to simulate limited Fields of View (FoV) (Fig. 1 (l)),

<sup>1</sup><https://github.com/gnns4hri/SocNavGym>

- 8) it allows limiting the FoV and sensor range of the agent (Fig. 1 (k)),
- 9) ORCA [33] and social force model [34] policies support for humans, with the option to consider or ignore robots,
- 10) optional and configurable noise for agents' sensory input,
- 11) manual control of the agent to collect LfD datasets.

#### A. Description of Gym features

SocNavGym is not dependent on a high-fidelity physics engine or any software framework to avoid additional overheads and delays while running simulations (feature 1). Previous works like SARL [10] and DSRNN [11] that use the likes of the proposed simulator only work with human entities in the environment. In our gym environment, the generated social scenarios can have different shapes and sizes across different episodic runs along with inanimate entities like plant, table, laptop and animate entities like human crowds of variable size (feature 2). In realistic social settings, human crowds and human-object interactions can form and break at any point in time, and our gym environment supports such dynamics (feature 3). This dynamic nature of interaction formation and dispersion is necessary to learn meaningful social context since the level of disturbance changes when an interaction undergoes a transition, that is, formation or dispersion.

Regarding the steering system of the robotic agents, previous works only support holonomic agents whose learnt policy cannot be directly transferred to non-holonomic agents. Thus we allow users to choose the agent they want to learn policies for (feature 4).

The reward functions used in previous works are arguably designed for collision avoidance rather than discomfort avoidance. To learn policies that truly account for social norms, we incorporated a discomfort estimator into our reward function that penalises agents for causing disturbance based on the context of the environment (feature 5). We further show a use case where we train a DuelingDQN agent with the proposed reward function and one of the commonly used in the literature and benchmark its behaviour using the evaluation metrics proposed in the guidelines paper [32], see Table II. The gaze of the human can be a good source of information to make navigation decisions. Therefore, for each human observed, SocNavGym provides a flag specifying whether the robotic agent is in their field of view (feature 7). The gym also enables users to mimic the constraints of their physical hardware on robots that might have limited field of view and range (feature 8). Humans in the gym can be controlled via ORCA or SFM policies (feature 9) and the users can choose whether the human behaviour should factor in the agent while calculating its policy.

To collect offline datasets, SocNavGym provides a sample script that allows the users to manually manipulate the agent in the room to collect (state, action, next state, reward) tuples (feature 11).

The observations take the form of dictionaries containing a key per entity type, the corresponding value being a list of such types of entities, namely "goal", "humans", "objects", "walls" and "relationships". All entities have basic properties including an identifier, 2D position coordinates, orientation, the radius of their bounding sphere, and their relative velocities -linear and angular. The observations returned by the environment can contain information about all the entities in the environment (fully observable) or only a subset of observable entities. This is configured by adjusting the range and field of view width of the robot, based on users' discretion (feature 8). Similarly, SocNavGym can be set up so that the observations are relative to the agent or a global frame of reference. In addition, the sensory information of specific entities can be corrupted with Gaussian noise by setting the distribution's parameters (feature 10). These variations of the observations can be configured via gym wrappers. If SocNavGym is configured to provide it, the "relationships" section of the observation space accounts for relational information (*e.g.*, humans talking, a human looking towards objects and human gaze, feature 7).

#### B. Evaluation metrics

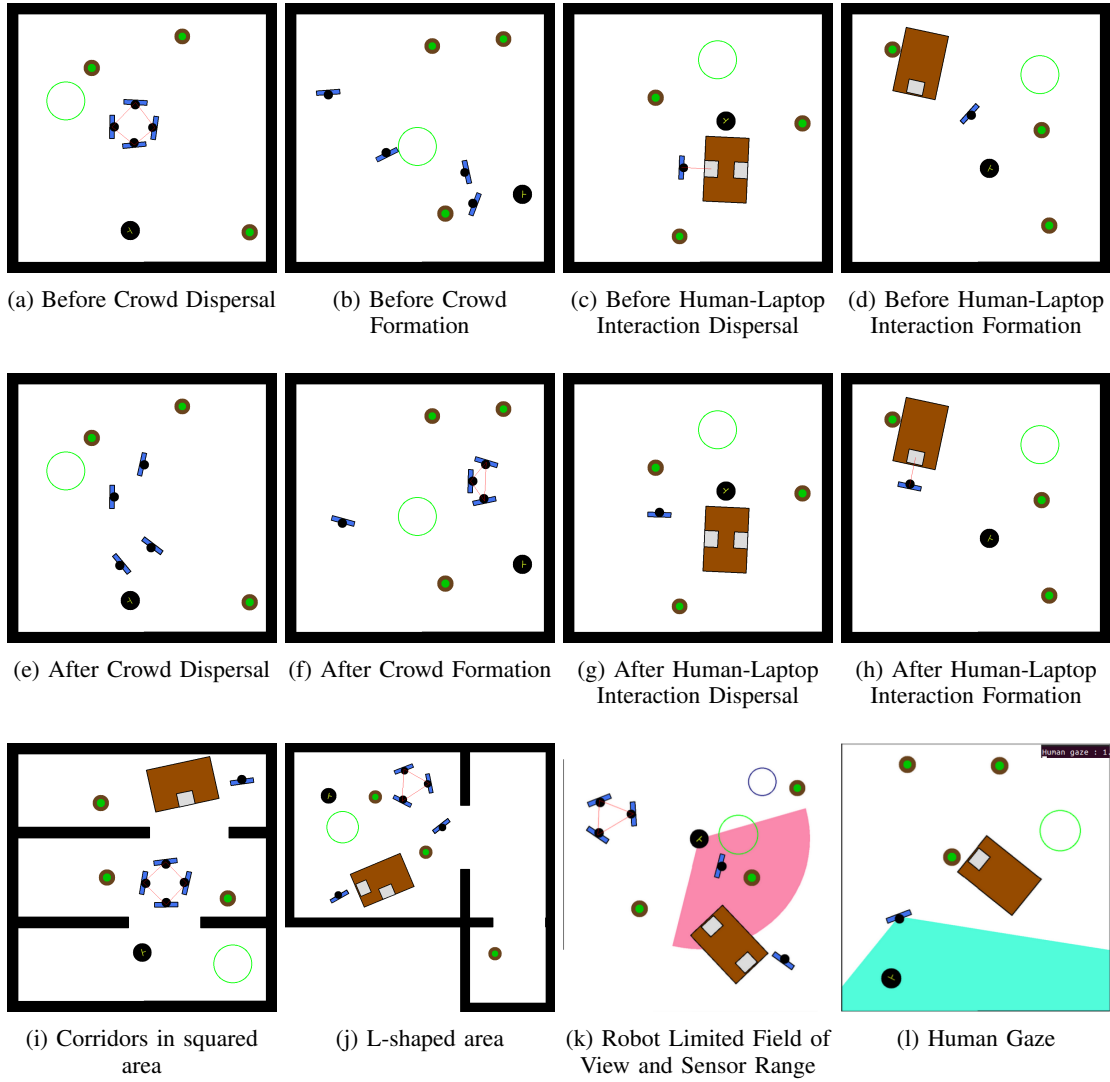
SocNavGym provides events and the evaluation metrics suggested in [32] via the `info` variable returned with each step of an episode (see Table I). They are designed to better understand agents' performance while training (feature 6). Along with these events and evaluation metrics, global metrics, corresponding to the execution of several episodes, can be obtained from the ones provided by the `info` variable. It also returns an adjacency matrix to identify the interacting entities in the environment at every instance of the simulation. These global metrics are accessible via helper functions to evaluate social agents in multiple runs.

#### C. Human trajectory generation

Simulated pedestrian trajectories can be generated with ORCA or a social force model. To get variable behaviours for each human, SocNavGym can be configured to use one of the two or to be chosen at random for every spawned pedestrian. The user can also choose for the robot to be ignored when calculating the simulated humans' actions. On spawning, the dynamic humans and crowds are given a goal location and, on reaching the goal, if the episode is not terminated, another goal location is sampled. SocNavGym can also be configured to create and dissolve groups, as well as human-human and human-object interactions. For crowd dispersion, each human or sub-group gets a different goal that they have to traverse or remain in place (static). Further, the user can change the FoV of the pedestrians to restrict its policy to only consider entities in their FoV.

#### D. Reward functions

SocNavGym allows making use of the reward function used by popular DRL-based SN approaches, (*e.g.*, [35], [10]), a piece-wise function that: a) gives a positive reward if the robot reaches the goal; b) provides a linear penalty if it gets



Legend						
Plant	Human	Robot	Robot Goal	Table with 2 laptops	Human-Laptop interaction	Human-Human interaction

(m) Entities' Labels

Fig. 1: Before (a-d) and after (e-h) scenario of dynamic formation and dispersal of a crowd and human-laptop interactions in the environment. Fig.i and Fig.j are square and L-shaped areas. Fig.k displays the sensor ranges of the robot and Fig.l shows the human gaze.

too close to a human; and c) rewards reducing the distance to the goal. In mathematical notation:

$$r(s_t, a_t) = \begin{cases} -1, & \text{if } d_{min}^t \leq r_{agent} + r_{entity} \\ (d_{min}^t - \delta_{disc}) \cdot \alpha \cdot dt, & \text{if } 0 < d_{min}^t < \delta_{disc} \\ 1.0, & \text{if } d_{goal}^t \leq \rho_{agent} \\ (d_{goal}^{t-1} - d_{goal}^t) \cdot \beta, & \text{otherwise} \end{cases} \quad (1)$$

where  $r(s_t, a_t)$  is the reward that the agent gets when performing action  $a_t$  in state  $s_t$  at time  $t$ ,  $d_{min}^t$  is the minimum acceptable distance between the robot and the closest human at time  $t$ ,  $r_{agent}$  and  $r_{entity}$  are the radius of the agent and an arbitrary entity in the environment respectively,  $\delta_{disc}$  (0.6m by default) is the threshold discomfort distance between the agent and the human,  $\alpha$  is a discomfort scaling factor,  $dt$  is the simulation timestep,  $\rho_{agent}$  is the goal radius,  $d_{goal}^t$  is the L2 distance to the goal position at time  $t$ , and  $\beta$  is a

Metric Name	Description
OUT_OF_MAP	True if the agent is out of the map
COLLISION_HUMAN & COLLISION_OBJECT & COLLISION_WALL	True if the agent has collided with a human OR object OR wall
COLLISION	True if the agent has collided with any entity
SUCCESS	True if the agent has reached the goal
TIMEOUT	True if the episode has terminated due to max episode length
FAILURE_TO_PROGRESS	The number of timesteps that the robot failed to reduce the distance to goal
STALLED_TIME	The number of timesteps that the robot's velocity is 0
TIME_TO_REACH_GOAL	Number of time steps taken by the robot to reach its goal
STL	Success weighted by time length
SPL	Success weighted by path length
PATH_LENGTH	Total path length covered by the robot
V_MIN & A_MIN & JERK_MIN	Minimum velocity OR acceleration OR jerk that the robot has achieved
V_AVG & A_AVG & JERK_AVG	Average velocity OR acceleration OR jerk of the robot
V_MAX & A_MAX & JERK_MAX	Maximum velocity OR acceleration OR jerk that the robot has achieved
TIME_TO_COLLISION	Minimum time to collision with a human at any point in time in the trajectory.
MINIMUM_DISTANCE_TO_HUMAN	Minimum distance to any human.
PERSONAL_SPACE_COMPLIANCE	Percentage of steps that the robot is not within the personal space (0.45m) of any human.
MINIMUM_OBSACLE_DISTANCE	Minimum distance to any object.
AVERAGE_OBSACLE_DISTANCE	Average distance to any object.
DISCOMFORT_DSRNN	The discomfort reward according to the DSRNN reward function
distance_reward	difference in agent's distance to goal in the previous and current time steps
sngnn_reward	The SNGNN reward that is $SNGNN\_value - 1. (DISCOMFORT\_SNGNN - 1)$
alive_reward	Alive reward that the agent received (meant to be used as a penalty)
closest_human_dist	Closest distance to a human
closest_obstacle_dist	Closest distance to an obstacle

TABLE I: Metrics in info dictionary




Fig. 2: Heatmaps of DSRNN-discomfort and SNGNNv2-discomfort for a social situation. Although the shape is similar (circular), the profile of the data-driven reward function is clearly non-linear. For visualisation purposes, we divided the room into an  $N \times N$  grid ( $N=100$ ) and placed the robot on each of these grid cells to obtain the discomfort values. We produce these values with respect to humans and interactions only so that a clean heatmap is obtained.

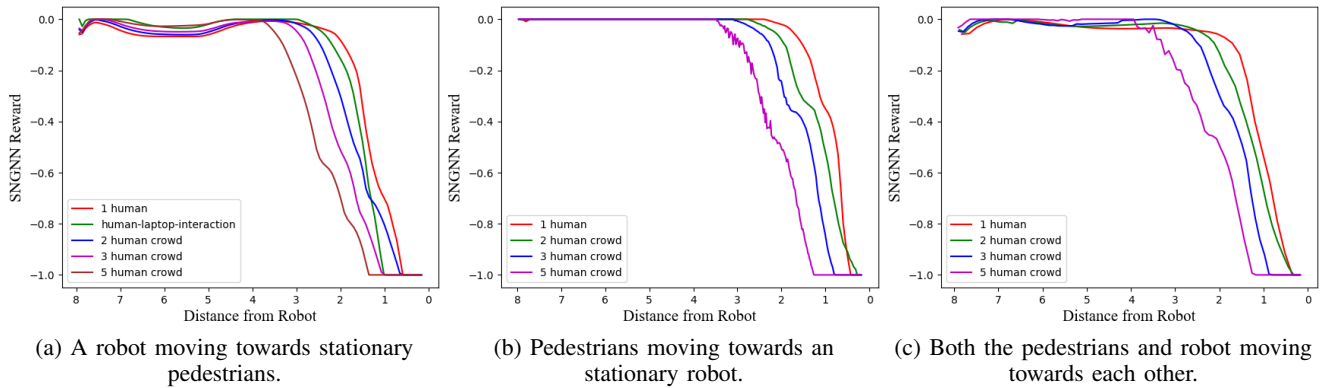


Fig. 3: SNGNN-DISCOMFORT when the agent (a) is moving towards the stationary entities, (b) is stationary and the entities are moving towards the agent, and (c) the entities are moving towards each other. SNGNN-DISCOMFORT considers the context as well as the presence or absence of different relations while evaluating the scenario.

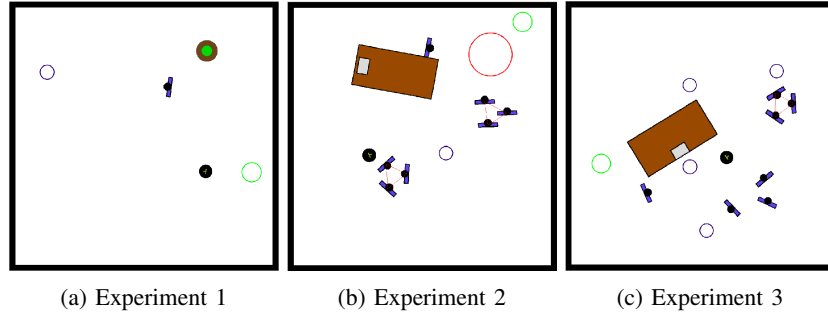


Fig. 4: Snapshot of the simulation environments used in the experiments.

scaling factor for the L2 distance.

The second piece of the reward function is a potential field around every entity in the room of radius  $\delta_{disc}$  within which the agent is penalised as its proximity from the entity decreases. This piece is the discomfort penalty that is used as a feedback to learn social compliant behaviour.

We argue that this reward function does not accurately account for human discomfort because it is agnostic to the context of the social situation. For instance, the threshold  $d_{min}^t$  is chosen arbitrarily and there are many variables that the function does not consider, such as velocities, orientations, or interactions. Another instance would be the agent getting a similar penalty when a human is approached from the back and front when, as studies have pointed out [36], the disturbance caused in both situations varies. Thus, we propose using a reward function that accounts for users' opinions based on the overall social context. As aforementioned, given the difficulty to model complex social preferences with hand-crafted equations, we propose the reward function in Eq. 2, which replaces the above discomfort factor with an SNGNN-v2 [19] based function. Leveraging a learned discomfort metric from human feedback arguably leads to better handling of social context:

$$r(s_t, a_t) = \begin{cases} -1, & \text{if } d_{min}^t \leq r_{agent} + r_{entity} \\ 1, & \text{if } d_{goal}^t \leq \rho_{agent} \\ (d_{goal}^{t-1} - d_{goal}^t) \cdot \beta + P(s_t), & \text{otherwise} \end{cases} \quad (2)$$

where  $P(s_t)$  is the SNGNN discomfort penalty, defined as:

$$P(s_t) = (SNGNN(s_t) - 1) \cdot \delta \quad (3)$$

$SNGNN(\cdot)$  being the function implemented by the SNGNN model and  $\delta$  being a scaling factor. SNGNN scores range from 1 (no disturbance) to 0 (extreme disturbance). Therefore, to use it as a penalty, 1 is subtracted from its output.

SNGNN-v2 [19] is trained to consider social phenomena. An example shown in Fig. 3 is how it accounts for the difference in the estimated discomfort in different situations given the robot-human distance. Similarly, Fig. 2 shows heatmaps of the discomfort values used in the two reward functions. DSRNN discomfort values are visibly homogeneous irrespective of the orientation of the human or the presence or absence of a crowd or interactions. This illustrates that

the DSRNN's reward function, although it is useful to study collision-free navigation, is arguably not the best to tackle the problem of Social Navigation.

#### IV. USE CASES

We integrate the Dueling-DQN agent with our environment to compare the performance of SNGNN based reward function with the one popularly used in the literature that we term DSRNN reward function. The agent neural network consists of 5 linear layers with 512, 256, 128, and 64 as the hidden units using LeakyRelu as the activation function. We used the Adam optimiser with a fixed learning rate of 0.001 and trained the agent separately with the two reward functions described above. The training lasted for 50K episodes with a maximum episode length of 200 discrete timesteps.

##### A. Social scenarios

The agent's observation space comprises of the position, orientation, and velocity of other entities in the agent's frame of reference. In each experiment, the position and orientation of all the entities' are initialised randomly, and the hyperparameters for the human's motion using ORCA or SFM are sampled from a Gaussian distribution for variable behaviour.

We train a Dueling DQN agent on three social scenarios that have increasing difficulty. All experiments were carried out with a non-holonomic agent. The following environment settings have been used to train the DRL agent:

- *Experiment 1*: In this experiment, we consider a 10m x 10m square room with 1 plant and 1 dynamic human (Fig 4.(a)).
- *Experiment 2*: A 10m x 10m square room with 1 table, 1 laptop placed randomly on the table, a static crowd of 3 humans, a dynamic crowd of 3 humans and one human that can dynamically form or break away an ongoing interaction with the laptop on the table during the episode (Fig 4.(b)).
- *Experiment 3*: A 10m x 10m square room with 1 table, 1 laptop placed randomly on the table, a static crowd of 3 humans, a dynamic crowd of 3 humans, and a dynamic human that can interact with the laptop. Humans interacting with the laptop and part of dynamic

	Exp 1 DSRNN	Exp 1 SNGNN	Exp 2 DSRNN	Exp 2 SNGNN	Exp 3 DSRNN	Exp 3 SNGNN
Discomfort SNGNN	-5.495	<b>-2.377</b>	-17.786	<b>-4.639</b>	-22.376	<b>-5.429</b>
Discomfort DSRNN	-0.2	<b>-0.143</b>	-0.481	<b>-0.403</b>	-1	<b>-0.341</b>
Personal Space Compliance	0.979	<b>0.996</b>	<b>0.997</b>	0.973	<b>0.998</b>	0.989
Closest Human Distance	<b>2.556</b>	2.528	1.456	<b>1.92</b>	1.376	<b>1.736</b>
Human Collision Rate	0.02	0.02	<b>0</b>	0.06	0.01	0.01
Wall Collision Rate	0.02	<b>0.01</b>	<b>0</b>	0.01	0.04	<b>0</b>
Object Collision Rate	0.07	<b>0.04</b>	<b>0.02</b>	0.06	0.12	<b>0</b>
Collision Rate	0.09	<b>0.06</b>	<b>0.02</b>	0.12	0.13	<b>0.01</b>
Minimum time to collision	186.511	<b>191.195</b>	172.24	<b>177.492</b>	174.56	<b>184.601</b>
Closest Obstacle Distance	1.039	<b>1.298</b>	0.806	<b>0.845</b>	0.751	<b>0.884</b>
Average Obstacle Distance	4.651	<b>4.686</b>	<b>4.449</b>	4.394	<b>4.458</b>	4.426
Timeout	<b>0.1</b>	0.26	0.5	<b>0.08</b>	0.53	<b>0.04</b>
Stalled Time	<b>18.61</b>	23.28	32.25	<b>13.72</b>	81.19	<b>0.37</b>
Time to Reach Goal	<b>90.55</b>	97.86	139.15	<b>79.77</b>	162.47	<b>67.14</b>
Failure to progress	<b>5.42</b>	12.27	29.06	<b>2.05</b>	<b>6.55</b>	7.84
Path Length	<b>5.558</b>	6.451	10.44	<b>4.575</b>	<b>6.288</b>	6.483
STL	<b>0.644</b>	0.607	0.388	<b>0.746</b>	0.204	<b>0.888</b>
SPL	<b>0.678</b>	0.624	0.393	<b>0.75</b>	0.27	<b>0.871</b>
Minimum velocity	0.03	0.042	0.03	0.053	0.006	0.082
Average velocity	0.08	0.084	0.082	0.09	0.052	0.098
Maximum velocity	0.1	0.1	0.1	0.1	0.1	0.1
Minimum acceleration	0.002	0.0	0	0.002	0	0
Average acceleration	0.023	0.031	0.064	0.011	0.018	0.012
Maximum acceleration	0.139	0.139	0.17	0.115	0.126	0.135
Minimum jerk	0.002	0.001	0	0.004	0.001	0
Average jerk	0.041	0.062	0.126	0.019	0.033	0.023
Maximum jerk	0.214	0.218	0.303	0.141	0.224	0.183

TABLE II: Performance of DuelingDQN when trained with the two reward functions. All the above metrics are averaged over 100 test episodes.

crowds are allowed to form and break their object or crowd interaction midst the episode ( Fig 4.(c))

More complex scenarios can be generated by modifying the variables in the configuration files.

### B. Results

The Dueling-DQN agent was trained using the two reward functions of Eq. 1 and Eq. 2 in the aforementioned social scenarios. Table II shows the performance of the trained agent across different evaluation metrics. Specifically, two different groups of metrics are considered: social compliance metrics and navigation metrics. Social compliance encompasses metrics such as SNGNN-discomfort (Eq. 2), DSRNN-discomfort (Eq. 1), Personal Space Compliance (PSC), closest distance to human and human collision rate. The remainder composes the navigation metrics. Minimum, average, and maximum velocity, acceleration, and jerk give additional information about the motion primitives of the agent. This second group is included to help users understand the agent's performance on traversing the goal location. As observed in Table II, according to the social metrics, the agent trained with the SNGNN-v2 reward function achieves higher social compliance when compared to the agent using the DSRNN reward. In addition, the performance of the SNGNN agent is also significantly better than the DSRNN agent in the navigation metrics in Experiments 2 & 3, which have a more complex social situation. This indicates that the agent can effectively learn to weigh the social context together with collision avoidance by means of this new discomfort penalty, which adds an additional benefit to the

use of our tool for the creation of new human-aware robot navigation algorithms.

## V. CONCLUSIONS AND FUTURE WORKS

In this work, we provide the HRI community with an easy-to-use social navigation environment capable to support the training of social navigation agents in different user-specified environment configurations that can vary in complexity, along with a social reward function. The experiments evidence that considering the social reward yields better social performance with little or no damage to the Success weighted by Time Length (STL). Although satisfactory, there are some limitations of the current work that we are working towards addressing in the near future. These include explicitly retraining SNGNN-v2 on holonomic systems and using more sophisticated navigation policies for humans.

### ACKNOWLEDGMENT

Most experiments were run in the Aston EPS Machine Learning Server, funded by the EPSRC Core Equipment Fund, Grant EP/V036106/1.

### 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. abs/2103.05668, p. 3583741, Feb. 2023. [Online]. Available: <https://dl.acm.org/doi/10.1145/3583741>
- [2] S. Liu, P. Chang, Z. Huang, N. Chakraborty, K. Hong, W. Liang, D. L. McPherson, J. Geng, and K. Driggs-Campbell, "Intention aware robot crowd navigation with attention-based interaction graph," 2022.



- [3] C. Chen, S. Hu, P. Nikdel, G. Mori, and M. Savva, "Relational Graph Learning for Crowd Navigation," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Las Vegas, NV, USA: IEEE, Oct. 2020, pp. 10 007–10 013.
- [4] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, and S. Savarese, "Social LSTM: Human Trajectory Prediction in Crowded Spaces," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. Las Vegas, NV, USA: IEEE, June 2016, pp. 961–971.
- [5] A. Gupta, J. Johnson, L. Fei-Fei, S. Savarese, and A. Alahi, "Social GAN: Socially Acceptable Trajectories with Generative Adversarial Networks," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. Salt Lake City, UT: IEEE, June 2018, pp. 2255–2264.
- [6] Y. Huang, H. Bi, Z. Li, T. Mao, and Z. Wang, "STGAT: Modeling Spatial-Temporal Interactions for Human Trajectory Prediction," in *IEEE/CVF International Conference on Computer Vision (ICCV)*. Seoul, Korea: IEEE, Oct. 2019, pp. 6271–6280.
- [7] S. Eiffert, H. Kong, N. Pirmarzdashti, and S. Sukkarieh, "Path planning in dynamic environments using generative rnns and monte carlo tree search," *CoRR*, vol. abs/2001.11597, 2020. [Online]. Available: <https://arxiv.org/abs/2001.11597>
- [8] M. L. Puterman, *Markov decision processes: discrete stochastic dynamic programming*, ser. Wiley series in probability and statistics. Hoboken, NJ: Wiley-Interscience, 2005.
- [9] P. Bachiller, D. Rodriguez-Criado, R. R. Jorvekar, P. Bustos, D. R. Faria, and L. J. Manso, "A graph neural network to model disruption in human-aware robot navigation," *Multimedia Tools and Applications*, vol. 81, no. 3, pp. 3277–3295, 2022.
- [10] C. Chen, Y. Liu, S. Kreiss, and A. Alahi, "Crowd-Robot Interaction: Crowd-Aware Robot Navigation With Attention-Based Deep Reinforcement Learning," in *International Conference on Robotics and Automation (ICRA)*. Montreal, QC, Canada: IEEE, May 2019, pp. 6015–6022. [Online]. Available: <https://ieeexplore.ieee.org/document/8794134/>
- [11] 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 (IROS)*. Vancouver, BC: IEEE, Sept. 2017, pp. 1343–1350. [Online]. Available: <http://ieeexplore.ieee.org/document/8202312/>
- [12] Z. Zhou, P. Zhu, Z. Zeng, J. Xiao, H. Lu, and Z. Zhou, "Robot navigation in a crowd by integrating deep reinforcement learning and online planning," *Applied Intelligence*, vol. 52, no. 13, pp. 15 600–15 616, Oct. 2022. [Online]. Available: <https://link.springer.com/10.1007/s10489-022-03191-2>
- [13] L. Qin, Z. Huang, C. Zhang, H. Guo, M. Ang, and D. Rus, "Deep Imitation Learning for Autonomous Navigation in Dynamic Pedestrian Environments," in *IEEE International Conference on Robotics and Automation (ICRA)*. Xi'an, China: IEEE, May 2021, pp. 4108–4115. [Online]. Available: <https://ieeexplore.ieee.org/document/9561220/>
- [14] J. P. de Vicente and A. Soto, "Deepsoconv: Social navigation by imitating human behaviors," *CoRR*, vol. abs/2107.09170, 2021. [Online]. Available: <https://arxiv.org/abs/2107.09170>
- [15] K. Narayanan, L. Posada, F. Hoffmann, and T. Bertram, "Acquisition of Behavioral Dynamics for Vision Based Mobile Robot Navigation from Demonstrations," *IFAC Proceedings Volumes*, vol. 46, no. 5, pp. 37–44, 2013. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S1474667015361942>
- [16] H. Kretzschmar, M. Spies, C. Sprunk, and W. Burgard, "Socially compliant mobile robot navigation via inverse reinforcement learning," *The International Journal of Robotics Research*, vol. 35, no. 11, pp. 1289–1307, 2016. [Online]. Available: <https://doi.org/10.1177/0278364915619772>
- [17] M. Fahad, Z. Chen, and Y. Guo, "Learning How Pedestrians Navigate: A Deep Inverse Reinforcement Learning Approach," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Madrid: IEEE, Oct. 2018, pp. 819–826. [Online]. Available: <https://ieeexplore.ieee.org/document/8593438/>
- [18] O. Saha and P. Dasgupta, "Improved reward estimation for efficient robot navigation using inverse reinforcement learning," in *NASA/ESA Conference on Adaptive Hardware and Systems (AHS)*. Pasadena, CA, USA: IEEE, July 2017, pp. 245–252. [Online]. Available: <http://ieeexplore.ieee.org/document/8046385/>
- [19] D. Rodriguez-Criado, P. Bachiller, and L. J. Manso, "Generation of human-aware navigation maps using graph neural networks," in *Artificial Intelligence XXXVIII*, M. Bramer and R. Ellis, Eds. Cham: Springer International Publishing, 2021, pp. 19–32.
- [20] I. Mordatch and P. Abbeel, "Emergence of Grounded Compositional Language in Multi-Agent Populations," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, no. 1, Apr. 2018. [Online]. Available: <https://ojs.aaai.org/index.php/AAAI/article/view/11492>
- [21] L. J. Manso, P. Nuñez, L. V. Calderita, D. R. Faria, and P. Bachiller, "SocNav1: A Dataset to Benchmark and Learn Social Navigation Conventions," *Data*, vol. 5, no. 1, p. 7, Jan. 2020. [Online]. Available: <https://www.mdpi.com/2306-5729/5/1/7>
- [22] R. Baghel, A. Kapoor, P. Bachiller, R. R. Jorvekar, D. Rodriguez-Criado, and L. J. Manso, "A toolkit to generate social navigation datasets," *CoRR*, vol. abs/2009.05345, 2020. [Online]. Available: <https://arxiv.org/abs/2009.05345>
- [23] Z. Wang, N. de Freitas, and M. Lanctot, "Dueling network architectures for deep reinforcement learning," *CoRR*, vol. abs/1511.06581, 2015. [Online]. Available: <http://arxiv.org/abs/1511.06581>
- [24] A. Dosovitskiy, G. Ros, F. Codevilla, A. M. López, and V. Koltun, "CARLA: an open urban driving simulator," *CoRR*, vol. abs/1711.03938, 2017. [Online]. Available: <http://arxiv.org/abs/1711.03938>
- [25] O. Saha and P. Dasgupta, "Improved reward estimation for efficient robot navigation using inverse reinforcement learning," in *NASA/ESA Conference on Adaptive Hardware and Systems (AHS)*. Pasadena, CA, USA: IEEE, July 2017, pp. 245–252. [Online]. Available: <http://ieeexplore.ieee.org/document/8046385/>
- [26] M. Campanella, S. Hoogendoorn, and W. Daamen, "The nomad model: Theory, developments and applications," *Transportation Research Procedia*, vol. 2, pp. 462–467, 2014, conference on Pedestrian and Evacuation Dynamics 2014.
- [27] S. Curtis, A. Best, and D. Manocha, "Menge: A modular framework for simulating crowd movement," *Collective Dynamics*, vol. 1, 03 2016.
- [28] N. Tsoi, A. Xiang, P. Yu, S. S. Sohn, G. Schwartz, S. Ramesh, M. Hussein, A. W. Gupta, M. Kapadia, and M. Vázquez, "Sean 2.0: Formalizing and generating social situations for robot navigation," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 11 047–11 054, 2022.
- [29] J. Holtz and J. Biswas, "SocialGym: A Framework for Benchmarking Social Robot Navigation," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Kyoto, Japan: IEEE, Oct. 2022, pp. 11 246–11 252. [Online]. Available: <https://ieeexplore.ieee.org/document/9982021/>
- [30] J. K. Haas, "A history of the unity game engine," Worcester Polytechnic Institute, Tech. Rep., 2014.
- [31] M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Ng, "Ros: an open-source robot operating system," in *IEEE International Conference on Robotics and Automation*, vol. 3, 01 2009.
- [32] A. Francis, C. Perez-D'Arpino, C. Li, F. Xia, A. Alahi, R. Alami, A. Bera, A. Biswas, J. Biswas, R. Chandra, H.-T. L. Chiang, M. Everett, S. Hall, J. Hart, J. P. How, H. Karnan, T.-W. E. Lee, L. J. Manso, R. Mirksy, S. Pirk, P. T. Singamaneni, P. Stone, A. V. Taylor, P. Trautman, N. Tsoi, M. Vazquez, X. Xiao, P. Xu, N. Yokoyama, A. Toshev, and R. Martín-Martín, "Principles and guidelines for evaluating social robot navigation algorithms," *arXiv preprint arXiv:2306.16740*, 2023.
- [33] 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*. Pasadena, California, USA: IEEE, 2008, pp. 1928–1935.
- [34] D. Helbing and P. Molnár, "Social force model for pedestrian dynamics," *Physical Review E*, vol. 51, no. 5, pp. 4282–4286, May 1995. [Online]. Available: <https://doi.org/10.1103/physreve.51.4282>
- [35] S. Liu, P. Chang, W. Liang, N. Chakraborty, and K. Driggs-Campbell, "Decentralized structural-rnn for robot crowd navigation with deep reinforcement learning," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 3517–3524.
- [36] E. T. Hall, R. L. Birdwhistell, B. Bock, P. Bohannon, A. R. Diebold Jr, M. Durbin, M. S. Edmonson, J. Fischer, D. Hymes, S. T. Kimball, et al., "Proxemics [and comments and replies]," *Current anthropology*, vol. 9, no. 2/3, pp. 83–108, 1968.