

# Socially Integrated Navigation: A Social Acting Robot with Deep Reinforcement Learning

Daniel Flögel<sup>1</sup>, Lars Fischer<sup>1</sup>, Thomas Rudolf<sup>1</sup>, Tobias Schürmann<sup>1</sup>, and Sören Hohmann<sup>2</sup>

**Abstract**—Mobile robots are being used on a large scale in various crowded situations and become part of our society. The socially acceptable navigation behavior of a mobile robot with individual human consideration is an essential requirement for scalable applications and human acceptance. Deep Reinforcement Learning (DRL) approaches are recently used to learn a robot’s navigation policy and to model the complex interactions between robots and humans. We propose to divide existing DRL-based navigation approaches based on the robot’s exhibited social behavior and distinguish between social collision avoidance with a lack of social behavior and socially aware approaches with explicit predefined social behavior. In addition, we propose a novel *socially integrated navigation* approach where the robot’s social behavior is adaptive and emerges from the interaction with humans. The formulation of our approach is derived from a sociological definition, which states that *social acting* is oriented toward the acting of others. The DRL policy is trained in an environment where other agents interact socially integrated and reward the robot’s behavior individually. The simulation results indicate that the proposed socially integrated navigation approach outperforms a socially aware approach in terms of ego navigation performance while significantly reducing the negative impact on all agents within the environment.

## I. INTRODUCTION

With the increasing development of smart cities, autonomous mobile robots are being integrated into our daily lives and deployed in various public pedestrian-rich environments [1]–[4]. As the robot performs its task, it navigates through a variety of scenarios, expecting to embody the respective human social behaviors [1, 5]. Humans follow social norms, e.g., keeping a distance from others, when navigating in crowded environments [6, 7]. However, they adapt their social behavior, e.g., the exact distance to others or the velocity, to the corresponding situation, and to that of the humans around them. The individual interpretation of appropriate social behavior is indirectly communicated to others through their movements and through social cues [7].

The navigation of a robot in crowded environments is generally referred to as social navigation and is performed by a social robot [8]–[10]. A socially intelligent robot has the same social abilities and skills as a human [11]. This field of research is, therefore, interdisciplinary and requires insights from engineering, human-robot cognition, psychology, design research, and sociology [9, 12]. This paper regards the social navigation problem from the sociological definition of *social acting* where acting is termed as social if it is

<sup>1</sup> are with FZI Research Center for Information Technology, Karlsruhe, Germany floegel@fzi.de

<sup>2</sup> is with the Institute of Control Systems at Karlsruhe Institute of Technology, Karlsruhe, Germany soeren.hohmann@kit.edu




Fig. 1. The proposed *socially integrated navigation* approach is adaptive to human behavior and preferences. It is based on a perspective change where the social behavior of the robot arises from its interaction with humans and is not predefined. The DRL policy has learned from the interaction behavior to consider the unknown personal space of each human.

related to the acting of others [13]. In addition, we consider the definition of *socially acceptable* behavior whereby the human-machine interaction is based on social norms, the individual nature of the human is taken into account, and the robot is able to adapt to human preferences as depicted in Fig. 1 [11].

Social navigation approaches are classified as local motion planning [6, 9] and can be divided into decoupled and coupled approaches [9]. Decoupled approaches first predict human motions and then plan collision-free trajectories [3]. However, due to the high uncertainty in human motions, this can lead to the freezing robot problem where no collision-free trajectory can be found [14]. In coupled approaches, the future crowd evolution is considered as joint sequential decision-making. Coupled approaches either assume a given structure of the problem as in game theory [15], the social force model (SFM) [16], and in optimal reciprocal collision avoidance (ORCA) [17] or assume insights about the principles as in Deep Reinforcement Learning (DRL) [18] or inverse reinforcement learning approaches [5, 19].

DRL approaches learn a navigation policy that maximizes the expected return over one episode [18]. Existing approaches either only show efficient navigation behavior while avoiding physical collisions [18, 20]–[22] or additionally exhibit a fixed social behavior, e.g., keeping a predefined distance to all humans [23]–[26] as illustrated in Fig. 2. The robot’s social behavior remains constant until retraining and every human in every scenario is treated equally regardless of how they behave. This contradicts socially acceptable behavior. These approaches exhibit no adaptive behavior, which is a requirement for social behavior from a sociological perspective [13]. In addition, a non-adaptive approach does

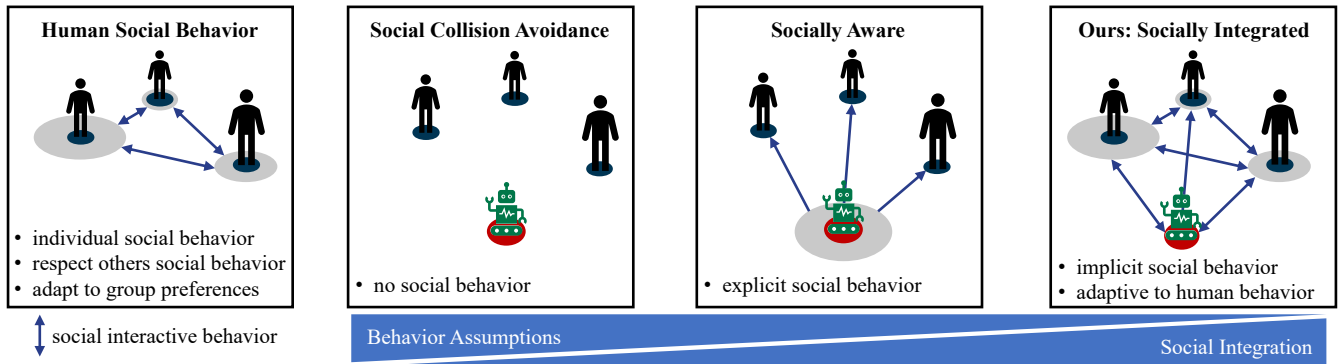


Fig. 2. We propose to distinguish between different DRL-based navigation approaches among pedestrians based on robot’s exhibited social interactive behavior. Social collision avoidance approaches consider only collision-free navigation. In socially aware approaches, a predetermined social behavior is projected onto all humans in the environment. Adaptive social behavior is considered in our socially integrated approaches, where humans’ individual social behavior is considered, and the robot’s behavior arises from an interaction with humans.

not scale to different scenarios and human behavior.

We address the above-mentioned drawbacks with a DRL-based *socially integrated navigation* approach, which adapts to the individual social behavior of humans in the current situation. The approach is derived from the sociological understanding of social acting and leads to socially acceptable robot behavior. We propose to rethink the navigation approaches and change the perspective of how the robot’s social behavior emerges. Moving away from an ego-robot perspective with fixed predefined social behavior toward mutual recognition as part of the crowd, where social behavior emerges through adaptation and the basic understanding of social norms. The socially integrated navigation formulation enables the robot to integrate itself into the crowd through its actions.

The main contributions of this paper are (i) a novel categorization of existing DRL-based social navigation approaches based on robot’s social behavior. (ii) We propose a DRL-based socially integrated navigation approach that empowers the robot to adapt its social behavior individually to human behavior and (iii) leverage principles of training a socially integrated policy from scratch. (iv) The evaluation with an impact analysis of the robot’s behavior in the crowd shows the improvements in social adaptive robot behavior.

## II. SOCIAL NAVIGATION

### A. Social Behavior in Crowded Environments

From a sociological perspective, the term *social* constitutes behavior that is related to the interaction between people and, based on Max Weber’s definition, is acting termed as social acting if it is related to the acting of other people and if the actions are oriented toward those of others [13]. In this context, the behavioral rules in the interaction are denoted as social norms [13], and the resulting social behavior emerges from reciprocal collision avoidance [6]. This implies that social behavior arises from an interaction between the behaviors of the individual participants and is not predetermined. In this regard, the interpersonal communication during the interaction is non-verbal through wordless cues, referred to as social cues, that are sent and received by humans [7]. These

social cues, such as facial expressions or motion patterns, control social interactions by expressing personal feelings and adapting individually to the reactions of others.

A human’s spatial management is described by the proxemics theory and was first proposed by Hall in [27]. Hall observed that people use spatial distances in various social and interpersonal interactions, whereby the distances vary depending on the environment and cultural factors. The spatial management is divided into the intimate, personal, social, and public zones, with increasing distance to the human. The dynamic and situation-dependent personal space is actively maintained by humans, and the intrusion of another human causes discomfort. Accordingly, humans try to maintain spaces for themselves that are preferred by themselves as well as by others [7].

For a robot to exhibit socially acceptable behavior, it must not only be guided by social norms but also view humans as individuals and adapt to their individual preferences [11]. The adaptation to context-specific social expectations, cultural norms, and individual preferences is a success factor for large-scale use of mobile robots in crowded environments [7, 11, 28, 29].

In summary, we conclude that a robot should have a basic understanding of social norms, but its social behavior should result from the individual interaction with a human through adaptation. In this way, a robot acts socially, has socially accepted behavior, and is applicable to large-scale use.

### B. Proposed Categorization for DRL Navigation Approaches

We propose to distinguish DRL-based navigation approaches in crowds based on the robot’s exhibited social behavior in human-machine interaction, as illustrated in Fig. 2. The social behavior is the basis for socially acceptable behavior as well as social acting and thus plays a fundamental role in the integration of robots into our society [1]. Recent reviews of social navigation have mainly focused on general methodological categorization and whether approaches demonstrate social behavior in principle [30, 31]. However, they do not directly distinguish for DRL-based approaches how they exhibit social behavior and how they

have learned social behavior. Thus, we distinguish between *Social Collision Avoidance*, *Socially Aware Navigation*, and propose the novel category *Socially Integrated Navigation* as described in the following.

1) *Social Collision Avoidance*: We refer to a robot navigation policy that is trained to efficiently reach the goal position while avoiding physical collisions in an environment with dynamic agents as social collision avoidance. Thereby, agents in the environment are referred to as humans but neither the robot nor the agents consider social norms or social behavior. The social aspect of these approaches arises from efficient and collision-free navigation next to humans. During the training procedure, the agent is rewarded for reaching the goal, penalized for colliding with other agents, and encouraged to efficient navigation [20]–[22, 32].

2) *Socially Aware Navigation*: We refer to a navigation policy that is, additionally to social collision avoidance, trained to respect explicit social behavior as socially aware navigation. Accordingly, the robot is rewarded for maintaining a predefined social behavior. This social behavior remains constant, and the robot interacts with every human in the same way, regardless of how the human behaves, as depicted in Fig. 2.

Training a robot to maintain a minimum distance from humans was first proposed with CADRL in [18] and augmented to an arbitrary number of humans in [23, 25] using a long short-term memory (LSTM) to encode the human-robot interaction. This pairwise interaction was augmented to attention-based crowd-robot interaction in [24], reformulated as a directed graph in [26] and augmented with an attention mechanism in [33, 34] to avoid performance degradation in large crowds. These approaches focus on efficient crowd-robot interaction modeling. Social behavior is restricted to maintaining a constant minimum distance from other humans, which can be seen as the robot’s personal zone.

Velocity-depending danger zones, which are differently shaped personal zones, were proposed in [35] and augmented with a hazardous area in [36]. Oriented bounding capsules were proposed in [10] to model non-circular shaped agents and velocity-depending risk-zones in moving direction. Emotional depending shaped personal spaces are proposed in [37] with a predefined set of comfort spaces. A risk-map-based approach with human position prediction is proposed in [38] to represent crowd interactions and geometric structures. These approaches consider variable personal zones but with predefined fixed assumptions from an ego-robot perspective.

A norm-inducing reward function is proposed in [6] to train an agent with overtaking, passing, and crossing capability. However, the agent is explicitly trained to do that in a predefined manner, e.g., passing on the right, and deviations from the predetermined social behavior are due to the tradeoff to time optimality and not due to adaptivity to human preferences. Hence, this is a predefined social behavior rather than an adaptation to human preferences.

All these approaches are based on predefined social behavior and show no adaptivity to human behavior. As a result, these approaches are not socially acceptable based

on the considered definitions [11, 13]. In addition, if the prior assumptions for social behavior do not apply, these approaches are not efficiently applicable.

3) *Socially Integrated Navigation*: We refer to a navigation policy as socially integrated, which is trained to exhibit implicit and adaptive social behavior to other agents’ behavior with individual consideration. Adaptation is not an imitation but rather a deviation from a personal understanding. In addition, the policy must be trained in an environment where other agents exhibit social behavior as well. There is no existing DRL approach that fulfills these minimum requirements toward a socially integrated navigation.

### III. PRELIMINARIES

Throughout this paper, a dynamic object in the environment is generally referred to as an agent, either a robot or a human, and its behavior is determined by a policy. Variables referred to the robot are indexed with  $x_0$ , and humans with  $x_i$  with  $i \in 1, \dots, N - 1$ . A scalar value is denoted by  $x$  and a vector by  $\mathbf{x}$ .

#### A. Problem Formulation

The navigation task of a robot toward a goal in a pedestrian-rich environment is a sequential decision-making problem and can be modeled as partially observable markov decision process (POMDP) [12, 18, 24]. The POMDP is described with a 7-tuple  $(\mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{O}, \mathcal{T}_0, R, \gamma)$ . We assume the state space  $\mathcal{S}$  and action space  $\mathcal{A}$  as continuous. The transition function  $\mathcal{T} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$  describes the probability transitioning from state  $\mathbf{s}_t \in \mathcal{S}$  to state  $\mathbf{s}_{t+1} \in \mathcal{S}$  for the given action  $\mathbf{a}_t \in \mathcal{A}$ . With each transition, an observation  $\mathbf{o}_t \in \mathcal{O}$  and a reward  $R : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$  is returned by the environment. The initial state distribution is denoted by  $\mathcal{T}_0$  while  $\gamma \in [0, 1)$  describes the discount factor. The POMDP serves as a framework for the problem formulation and is solved with a DRL approach [39]. The environment contains  $N$  agents with one robot ( $i = 0$ ) and  $N - 1$  humans ( $i \in 1, \dots, N - 1$ ). In single robot navigation, only the robot’s navigation policy  $\pi_0$  is trained. The human policies  $\pi_i$  are unknown to the robot and correspond to an existing behavioral model. Every agent is completely described with a state  $\mathbf{s}_{i,t} = [\mathbf{s}_{i,t}^o, \mathbf{s}_{i,t}^h]$  at any given time  $t$ , where the state is separable into an observable  $\mathbf{s}_{i,t}^o$  and unobservable  $\mathbf{s}_{i,t}^h$  part:

$$\mathbf{s}_i^o = [\mathbf{p}, \mathbf{v}, r] \in \mathbb{R}^5 \quad (1)$$

$$\mathbf{s}_i^h = [\mathbf{p}_g, v_{\text{pref}}, \psi_{\text{pref}}, r_{\text{prox}}] \in \mathbb{R}^5 \quad (2)$$

The observable state is known to everyone and composed of the position  $\mathbf{p} = [p_x, p_y]$ , velocity  $\mathbf{v} = [v_x, v_y]$ , and radius  $r$ . We assume agents as circular objects where the radius determines the area which is occupied by this agent. This assumption is valid since we consider an environment with only humans and a robot with a human-like shape. For agents with other shapes, oriented bounding capsules are proposed in [10]. The unobservable state is only known to the agent itself and composed of goal position  $\mathbf{p}_g = [p_{gx}, p_{gy}]$ , preferred velocity  $v_{\text{pref}}$ , preferred orientation  $\psi_{\text{pref}}$ , and a proxemic radius  $r_{\text{prox}}$ . The proxemic radius represents the

agent’s personal space according to Hall’s proxemic theory. The world state  $\mathbf{s}_t = [s_{0,t}, \dots, s_{N,t}]$  represents the current situation in the environment.

We regard the POMDP as observable given sufficient partial history, whereby the hidden states of an agent can be inferred from past observations. In every timestep  $t$ , the robot observes the environment receiving  $\mathbf{o}_t$  and takes an action  $\mathbf{a}_t = \pi(\mathbf{a}_t|\mathbf{o}_t)$ . The robot kinematic is assumed to be a unicycle model and the action  $\mathbf{a}_t = [v_t, \Delta\theta_t]$  is the velocity and heading change, respectively. The trajectory  $\tau$  of one episode is the sequence of states, observations, actions, and rewards within the terminal time  $T$ . The return of one episode  $\mathcal{R}(\tau) = \sum_{t=0}^T \gamma^t R_t$  is the accumulated discounted reward  $R_t$  and the central objective is to learn the optimal robot policy  $\pi^*$  which maximizes the expected return:

$$\mathcal{T}(\tau|\pi) = \mathcal{T}(s_0) \prod_{t=0}^T \mathcal{T}(s_{t+1}|s_t, \mathbf{a}_t) \pi(\mathbf{a}_t|\mathbf{o}_t) \quad (3)$$

$$\mathbb{E}_{\tau \sim \pi} [\mathcal{R}(\tau)] = \int_{\tau} \mathcal{T}(\tau|\pi) \mathcal{R}(\tau) \quad (4)$$

$$\pi^*(\mathbf{a}|\mathbf{o}) = \arg \max_{\pi} \mathbb{E}_{\tau \sim \pi} [\mathcal{R}(\tau)] \quad (5)$$

Considering a stochastic environment,  $\mathcal{T}(\tau|\pi)$  is the probability of a trajectory starting in  $s_0$  with the probability  $\mathcal{T}_0$ .

#### IV. APPROACH

This section presents our DRL-based socially integrated navigation approach, which is based on social acting and learns to adapt to the social behavior of humans. We first describe how we incorporated the social integration formulation within a DRL framework, and subsequently, we impose minimum requirements on the training process.

##### A. Adapting Social Behavior

Recall that the objective of the DRL training process is to find a policy that maximizes the expected accumulated discounted reward over one episode. The reward system is decisive for what the policy has learned from the observation, whereas the observation must be sufficient to learn the desired behavior. In the environment, humans have their own objectives and evaluate the current situation and the behavior of others based on their personal preferences and objectives. This assessment is then transferred to other people in the environment through social cues.

This leads to an environment where multiple reward systems exist, as depicted in Fig. 3. We interpret the social cues as a reward a human will give the robot. We assume an integrated society where everyone is perceived as part of the group, and therefore, we assume that people treat the robot equally to humans. This individual human reward evaluates the robot’s behavior from the individual human perspective, and all rewards together evaluate the robot’s behavior in the group. Since not every human in the environment gives social cues, e.g., because the distance is too large or he is not directly involved in the interaction, only the people within a certain radius  $r_{SI}$  around the robot are considered as illustrated in Fig. 3. In addition, we consider the adaptation




Fig. 3. Proposed distributed rewards in the environment. The robot has only one navigation reward and receives the reward for the social behavior of humans within the social integration radius  $r_{SI}$ .

of social behavior as a local process. From this principle, we derive the socially adaptive reward

$$R_{SA} = \frac{1}{M} \sum_{i=1}^M \lambda_i R_i \quad (6)$$

that the robot receives at each time step. With the reward  $R_i$  of human  $i$ , a scaling factor  $\lambda_i$ , and the number of humans  $M$  within the social integration radius  $r_{SI}$  around the center of the robot. If the human  $i$  is within the social integration radius of the robot  $d_{0i} < r_{SI}$ , the human reward is given at every time step by

$$R_i = -R_v \cdot |v_{i,t} - v_{0,t}| - R_{prox}. \quad (7)$$

The reward comprises velocity deviations scaled with  $R_v$  and rewards violations of humans’ personal space  $d_{i0} < r_{i,prox}$  with  $R_{prox}$ . We observed that using sparse reward for proxemic violations leads to better results and generalization than dense reward. The scaling factor allows the robot to interpret and evaluate the received reward and to pay more attention to certain humans. For example, more focus can be paid to people who are close or further away from the robot. The social integration radius is decisive for how granularly local group situations are regarded. Note that being socially integrated does not mean imitating others. The robot forms its understanding and interpretation of the current situation from its perspective and adapts to the group behavior, taking people individually into account. In addition to the social adaptive reward, the robot is rewarded for efficient navigation with  $R_{Nav}$  which leads to the reward at time step  $t$

$$R_t = R_{Nav} + R_{SA}. \quad (8)$$

The navigation reward is responsible for teaching the robot efficient navigation toward the goal while avoiding collisions.

$$R_{Nav} = \begin{cases} +R_g & \text{if } \mathbf{p}_t = \mathbf{p}_g \\ -R_c & \text{if } d_{0i} \leq 0 \\ -R_{time} & \text{if timeout} \\ +R_{gd,1} \cdot |\Delta d_g| & \text{if } \Delta d_g > 0 \\ -R_{gd,2} \cdot |\Delta d_g| & \text{if } \Delta d_g < 0 \end{cases} \quad (9)$$

We encourage to make steps toward the goal with  $R_{gd,1}$  with  $\Delta d_g = d_{g,t} - d_{g,t-1}$  and reaching the goal with  $R_g$ . Contrarily, we penalize the robot for collisions with other agents with  $R_c$ , running into timeouts with  $R_{time}$ , and making

steps from goal away with  $R_{gd,2}$ . The dense navigation rewards support fast and stable training.

A key requirement to adapt to the social behavior of others is that the others follow their social norms, and the robot can observe their behavior. This means that the social norms must be observable based on the observation the agent does, and the social norms must be observable based on the interaction with the other agent. This implies that the agent must consider the temporal crowd evolution and spatial relationships between the agents.

The observation is divided into the robot observation  $\mathbf{o}_{0,t}$  and human observation  $\mathbf{o}_{i,t} = [\bar{\mathbf{o}}_i, \hat{\mathbf{o}}_{i,t}]$ . Based on human's observable states, the robot observes the human  $i$  at time  $t$ :

$$\mathbf{o}_{0,t} = [d_g \quad \Delta \mathbf{p}_g \quad \theta \quad v_{\text{pref}} \quad r_0] \quad (10)$$

$$\bar{\mathbf{o}}_i = [r_i \quad r_i + r_0] \quad (11)$$

$$\hat{\mathbf{o}}_{i,t} = [d_t \quad \Delta \mathbf{p}_t \quad \Delta \mathbf{v}_t] \quad (12)$$

The robot observes its state based on the relative position to the goal  $\Delta \mathbf{p}_g = \mathbf{p}_g - \mathbf{p}_t$ , the direct distance to the goal  $d_g = \|\Delta \mathbf{p}_g\|_2$ , the heading  $\theta$ , the preferred velocity  $v_{\text{pref}}$ , and the personal radius  $r$ . We distinguish the human observation in a constant part  $\bar{\mathbf{o}}_i$  which does not change over time, and a time-varying part  $\hat{\mathbf{o}}_{i,t}$  which varies throughout the steps. The constant part contains the human radius  $r_i$  and the combined radius  $r_i + r_0$ . The temporal observation  $\hat{\mathbf{o}}_{i,t}$  contains the relative position  $\Delta \mathbf{p}_t = \mathbf{p}_{0,t} - \mathbf{p}_{i,t}$ , direct distance  $d_t = \|\Delta \mathbf{p}_t\|_2$ , and relative velocity  $\Delta \mathbf{v}_t = \mathbf{v}_{0,t} - \mathbf{v}_{i,t}$  between robot and human. To consider the partial history, we concatenate the last  $k$  temporal observations to the aggregated observation  $\mathbf{o}_{i,t}$  of human  $i$  at timestep  $t$

$$\mathbf{o}_{i,t} = [\bar{\mathbf{o}}_{i,t} \quad \hat{\mathbf{o}}_{i,t} \quad \hat{\mathbf{o}}_{i,t-1} \quad \cdots \quad \hat{\mathbf{o}}_{i,t-k}]. \quad (13)$$

In a human-aware perspective, the joint observation

$$\mathbf{o}_t = \begin{bmatrix} \mathbf{o}_{0,t}, \mathbf{o}_{1,t} \\ \cdots \\ \mathbf{o}_{0,t}, \mathbf{o}_{N-1,t} \end{bmatrix} \quad (14)$$

encodes the pairwise robot-human interaction capacity and temporal evolution.

### B. Training a Socially Integrated Agent

The formulation of our socially integrated navigation approach based on social acting refers to the behavior of others and adapting to their behavior. In a real-world setting, the robot is used in an environment where humans respect the social behavior of others and have learned to adapt to their behavior, and they interact socially integrated. To ensure realistic and representative training, we impose the minimum requirement to the training environment that humans act according to their own social behavior and that of other humans. This implies that they actively avoid violating their own preferences and those of others. In addition, the social behavior of humans must correspond to the social norms to which the robot should adapt in the real world. Otherwise, the robot can't infer social behavior from social cues and observations, and thus, it is not possible to train a socially integrated robot.




Fig. 4. Proposed perspective of the various agents among each other in the environment. The gray-shaded personal areas are only visible to the respective agent. Humans act in the environment to maintain their personal space and that of others.

We follow the proposal to use the ORCA model to simulate human interactive behavior [25, 26, 38, 40] since it is more stable than SFM-based approaches in densely crowded scenarios [30]. In a nutshell, ORCA computes a collision-free velocity for the ego robot within a time horizon, assuming that other agents in the environment follow the same policy, and thus ORCA does not take into account direct social behavior. A cooperation coefficient is used to split the collision avoidance effort among all agents. As the robot is a technical machine, it has no personal space, only humans. In an environment with one robot and  $N - 1$  humans, there are three basic perspectives: robot-human, human-robot, and human-human. The key difference in these perspectives is how others are perceived and how they interact with others. Fig. 4 illustrates the perspectives for the case where each agent has its personal space and humans act socially integrated. Recall, according to the observable and unobservable state in (1) and (2) of an agent, the human proxemic radius  $r_{\text{prox}}$  is not within the observation. Only the radius  $r_i$  and position of each agent can be observed. This leads the robot to observe the distance  $d_{0i}$  between the robot and human  $i$  since the robot can not observe the hidden states. To allow the ORCA models to behave appropriately according to individual social behavior, we give these models general knowledge about the entire environment. Therefore, humans aim to maintain  $d_{i0} \geq 0$  when interacting with the robot. In human-human interaction, they react according to  $d_{ij} \geq 0$  and thus respect each proxemic radius. This implies that the robot can deduce individual human behavior from human-robot and human-human interaction observations.

The navigation policy is trained with the proximal policy optimization (PPO) [41] algorithm by simultaneously learning a critic network and a policy. We follow [25] to use a LSTM as a feature extractor to handle a varying number of humans but use separated policy and critic networks. A comparison across 15 seeds showed that a shared feature extractor architecture, as depicted in Fig. 5, converges faster to a mean success rate of 1 with higher reward and better navigation metrics. Recent approaches proposed to apply imitation learning with expert trajectories in advance of the DRL training process to initialize the policy for stable and converging training [10, 22, 24]–[26, 33]. However, imitation learning with expert knowledge introduces a bias through the expert behavior. Thus, we propose to reward the robot's velocity to avoid imitation learning. This ensures that the




Fig. 5. Proposed architecture with shared features extractor and separated policy and critic networks. The last hidden state of the LSTM is used as input for fully connected layers.

policy can be trained from scratch and has no bias caused by expert trajectories. In our socially integrated formulation, a velocity reward can either be incorporated into the navigation reward  $R_{Nav}$  with adding  $-R_v \cdot |v_{i,t} - v_{i,pref}|$  to train the robot having one preferred velocity or into the  $R_{SA}$  to train the robot adapting to the velocity of others as denoted in (7).

## V. EVALUATION

### A. Experimental Setup

We consider three different scenarios for evaluation. One passing scenario representing a pedestrian zone, depicted in Fig. 6 column 1 and 2, and two circle crossing scenarios representing densely crowded environments, depicted in Fig. 6 column 3 and 4 on the right side. Throughout one episode, we keep the hidden state  $s_i^h$  of all agents constant, and all agents have the same radius  $r_i = 0.3$ . The circle crossing scenario includes 8 agents with uniformly sampled start and end positions. We distinguish between a homogeneous scenario, where each human has the same random proxemic radius, and a heterogeneous scenario, where each human has its random individual proxemic radius. In both circle crossing scenarios, the proxemic radius  $r_{prox}$  is uniformly sampled from  $[0, 0.8]$ , according to the intimate and personal zone in [27], and the human velocity is uniformly sampled between  $[0.5, 1]$ . Humans are controlled by an ORCA policy and additionally respect the social behavior of others according to the training procedure for socially integrated agents. The cooperation coefficient is uniformly sampled from  $[0.3, 0.7]$  whereas 0.5 means that the agent will apply half the effort to avoid a collision within the time horizon of 5s. Since we aim to integrate the robot into the human group and the society, we do not evaluate our approach in scenarios where the robot is invisible to humans as done in [24, 36].

To conduct the experiments and train the DRL policy, we developed a customized gymnasium environment [42] inspired by [25] and incorporated the DRL framework *Stables Baselines3* [43]. We compare our socially integrated agent against a socially aware agent with explicit social behavior considering spatial observation ( $k = 0$ ) according to [23]–[26, 33] and temporal observation ( $k = 15$ ). The socially aware agent has the same policy architecture and navigation reward  $R_{Nav}$  but the robot’s social behavior emerges through an ego perspective of (7) with  $r_{0,prox} = 0.2$  to maintain a minimum distance to all humans and a predefined velocity.

TABLE I  
PPO HYPERPARAMETERS AND REWARDS

Hyperparameter	Value	Reward	Value
optimizer	Adam	$R_g$	4
learning rate	$3 \times 10^{-4}$	$R_c$	4
environment steps	2048	$R_{time}$	4
clip range	0.2	$R_{gd,1}$	0.1
number of epochs	2	$R_{gd,2}$	0.2
batch size	64	$R_v$	0.052
discount factor $\gamma$	0.99	$R_{prox}$	1.1
activation functions	ReLU		

### B. Results

1) *Avoiding Imitation Learning*: We trained the socially integrated and socially aware policy over 15 seeds, using the hyperparameters and rewards from Table I, and observed convergence after 1.8M and 2M steps, respectively, across all seeds. The temporal observation includes the human motions of the last 3s, leading to  $k = 15$  with a discrete step size of  $\Delta t = 0.2$  s. To avoid imitation learning with expert trajectories, we investigate the impact of the introduced velocity reward on the training behavior. We train the policy with ( $R_v = 0.052$ ) and without ( $R_v = 0$ ) velocity reward in a circle crossing scenario over 15 seeds, each with a maximum of 10M steps. The navigation performance of the trained policies reveals that using a velocity reward in training leads to a navigation policy with task completion in all 15 seeds, whereas without the velocity reward, only 1 trained policy could solve the navigation task.

2) *Adapting Social Behavior*: To prove the adaptivity of our approach to individual human behavior and the impact on others, we consider the same passing scenario twice as depicted in Fig. 6 on the left half. In column 1, the approaching other agents have an invisible personal space of  $r_{prox} = 0.8$ , and in column 2, they have no personal space. The socially integrated robot successfully adapts the distance to the approaching agents depending on the humans’ personal space, whereas the socially aware agent violates the other humans’ personal space while moving straight toward the goal. Although the human acts to maintain his personal space. In contrast, the socially integrated agent accelerates so that he passes the green agent earlier when he is already interacting with the purple agent. If the passing agents have no personal space, both robots perform similarly. In the heterogeneous scenario in column 4 of Fig. 6, the socially aware robot violates the personal space of the yellow agent and also pushes him back as the yellow agent acts to maintain its personal space. In contrast, the socially integrated robot passes the other agents on the other side and thus avoids the conflict.

The ablation study, based on the heterogeneous (HE) and homogeneous (HO) circle crossing scenarios, is aggregated in Table II for a socially integrated and a socially aware robot with history ( $k=15$ ) and without ( $k=0$ ). The two ratios indicate how much more or less time and distance is required compared to the ideal value. The human return value aggregates (7). The socially aware robots cause more collisions, run into more timeouts, and cause significantly more proxemic violations. A socially aware robot with

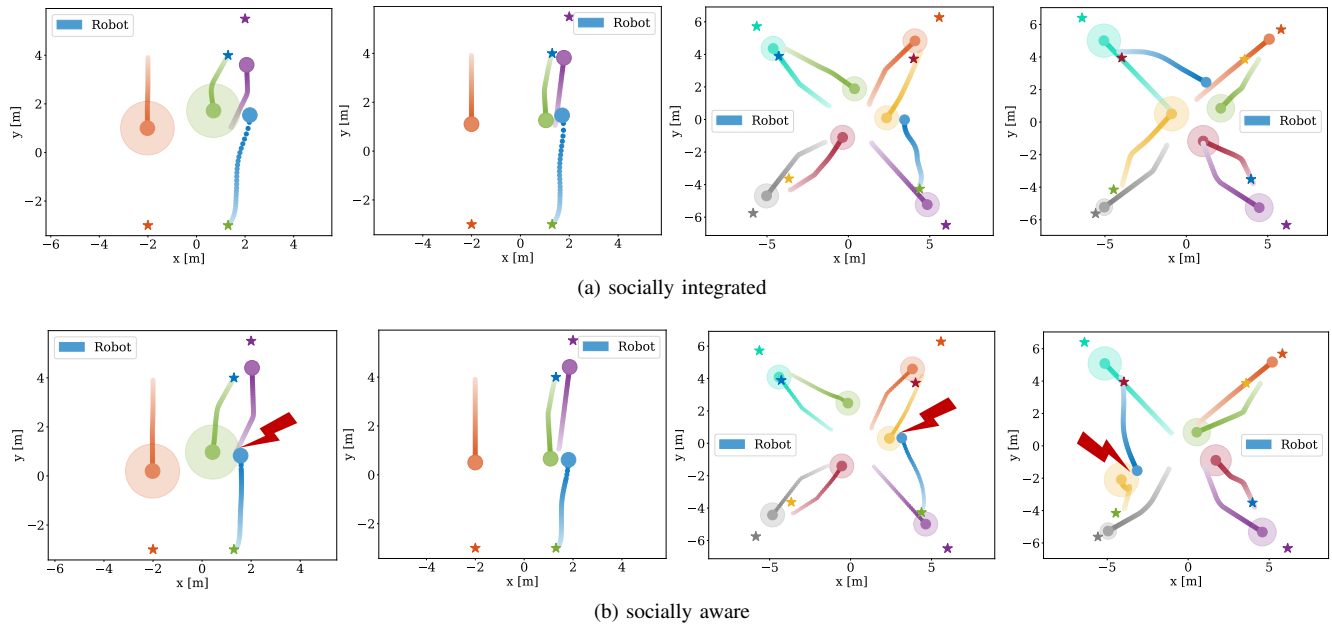


Fig. 6. Comparison of our proposed socially integrated navigation approach with adaptive social behavior in (a) to a socially aware navigation approach with fixed social behavior in (b). A passing scenario is depicted in column 1 and 2 on the left side and a homogenous and heterogenous circle crossing scenario in column 3 and 4 to the right. Humans’ personal space, depicted with faded circles around an agent, is invisible to the robot and uniformly sampled. Violations of humans’ personal space are marked with a red spark.

history has improved ego navigation performance but more negative impact on humans, measured by more negative human return, more proxemic violations, and worse human time ratio. In comparison, the socially integrated robot has better ego navigation performance, causes less human discomfort, reflected in human return and proxemic violations, and leads humans to slightly better distance and time ratios.

### C. Discussion

Through the extension of the ORCA model with individual proxemic radius and the environment design that humans respect the proxemics of others, the simulated human social behavior is closer to the reality than just using the ORCA model as used in previous approaches. Therefore, the environment is well suited to investigate the effect of the robot’s exhibited social behavior. For the socially aware navigation, the results clearly show a trade-off between an improvement in robot navigation performance and causing negative effects on other agents in the environment. At the same time, the socially aware agents do not react to people’s conscious actions and push them back, for example. In contrast, the socially integrated approach does not have this trade-off. The perspective change leads to a significantly reduced negative impact on others and at the same time to better robot navigation performance. In addition, due to its ability to adapt to the individual behavior of others, the socially integrated robot can react more consciously to the reactions of others and thus make more far-sighted decisions. However, if others are not cooperative and do not interact, the robot can not adapt, and proxemic violations can still occur. Overall, a socially integrated robot outperforms a socially aware robot.

## VI. CONCLUSIONS

In this paper, a distinction between social collision avoidance, socially aware navigation, and socially integrated navigation approaches is proposed with regard to the exhibited social behavior of the robot. The proposed socially integrated navigation approach is derived from sociological definitions and an understanding of human behavior in crowded environments. The results show that the proposed perspective change from predefined explicit social behavior to the adaption to individual human behavior leads to better robot navigation while reducing the negative impact on all agents involved. Consequently, robots should have the same social understanding and decision-making as humans and thus act in a socially integrated manner. Moreover, our approach is applicable in a scalable manner through adaptive behavior. Future work includes an evaluation in the real world.

## REFERENCES

- [1] F. Cavallo, F. Semeraro, L. Fiorini, G. Magyar, P. Sinčák, and P. Dario, “Emotion modelling for social robotics applications: A review,” *Journal of Bionic Engineering*, vol. 15, no. 2, pp. 185–203, 2018.
- [2] E. Alao and P. Martinet, “Uncertainty-aware navigation in crowded environment,” in *17th International Conference on Control, Automation, Robotics and Vision (ICARCV)*. IEEE, 2022, pp. 293–298.
- [3] B. Varga, D. Yang, M. Martin, and S. Hohmann, “Cooperative decision-making in shared spaces: Making urban traffic safer through human-machine cooperation,” in *IEEE 21st Jubilee International Symposium on Intelligent Systems and Informatics (SISY)*, 2023.
- [4] T. Genevois, A. Spalanzani, and C. Laugier, “Interaction-aware predictive collision detector for human-aware collision avoidance,” in *2023 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2023, pp. 1–7.
- [5] D. Gonon and A. Billard, “Inverse reinforcement learning of pedestrian–robot coordination,” *IEEE Robotics and Automation Letters*, vol. 8, no. 8, pp. 4815–4822, 2023.
- [6] C. Yufan, E. Michael, L. Miao, and P. H. Jonathan, *Socially Aware Motion Planning with Deep Reinforcement Learning*. Piscataway, NJ: IEEE, 2017.

TABLE II

ABLATION STUDY: OUR SOCIALLY INTEGRATED AND A SOCIALLY AWARE (SA) AGENT WITH AND WITHOUT HISTORY IN 100 RANDOM EPISODES.

Scenario	Method	Col.	Time-out	# Proxemic violations	Distance ratio ( $\mu \pm \sigma$ )		Time ratio ( $\mu \pm \sigma$ )		Human Return ( $\mu \pm \sigma$ )
					Robot	Human	Robot	Human	
HO	SA (k=0)	2	3	177	1.24 $\pm$ 0.24	1.03 $\pm$ 0.04	1.06 $\pm$ 0.23	1.10 $\pm$ 0.13	-1.50 $\pm$ 2.36
	SA (k=15)	1	0	554	1.11 $\pm$ 0.06	1.03 $\pm$ 0.06	1.05 $\pm$ 0.14	1.12 $\pm$ 0.16	-1.76 $\pm$ 3.59
	Ours	<b>0</b>	<b>0</b>	<b>41</b>	1.14 $\pm$ 0.04	<b>1.03 <math>\pm</math> 0.03</b>	<b>1.00 <math>\pm</math> 0.11</b>	<b>1.10 <math>\pm</math> 0.13</b>	<b>-1.29 <math>\pm</math> 1.24</b>
HE	SA (k=0)	2	2	197	1.25 $\pm$ 0.27	1.03 $\pm$ 0.03	1.08 $\pm$ 0.22	1.10 $\pm$ 0.11	-1.63 $\pm$ 2.06
	SA (k=15)	2	0	632	1.11 $\pm$ 0.05	1.03 $\pm$ 0.03	1.05 $\pm$ 0.10	1.10 $\pm$ 0.11	-1.93 $\pm$ 3.34
	Ours	<b>0</b>	<b>0</b>	<b>34</b>	1.14 $\pm$ 0.03	<b>1.02 <math>\pm</math> 0.03</b>	<b>1.01 <math>\pm</math> 0.09</b>	<b>1.10 <math>\pm</math> 0.10</b>	<b>-1.38 <math>\pm</math> 1.04</b>

- [7] J. Rios-Martinez, A. Spalanzani, and C. Laugier, "From proxemics theory to socially-aware navigation: A survey," *International Journal of Social Robotics*, vol. 7, no. 2, pp. 137–153, 2015.
- [8] T. Fong, I. Nourbakhsh, and K. Dautenhahn, "A survey of socially interactive robots," *Robotics and Autonomous Systems*, 2003.
- [9] C. Mavrogiannis, F. Baldini, A. Wang, D. Zhao, P. Trautman, A. Steinfield, and J. Oh, "Core challenges of social robot navigation: A survey," 09.03.2021.
- [10] K. Zhu, B. Li, W. Zhe, and T. Zhang, "Collision avoidance among dense heterogeneous agents using deep reinforcement learning," *IEEE Robotics and Automation Letters*, vol. 8, no. 1, pp. 57–64, 2023.
- [11] K. Dautenhahn, "Socially intelligent robots: dimensions of human-robot interaction," *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, no. 1480, pp. 679–704, 2007.
- [12] K. Zhu and T. Zhang, "Deep reinforcement learning based mobile robot navigation: A review," *Tsinghua Science and Technology*, vol. 26, no. 5, pp. 674–691, 2021.
- [13] H. Korte and B. Schäfers, Eds., *Einführung in Hauptbegriffe der Soziologie*, 9th ed., ser. Einführungskurs Soziologie. Wiesbaden: Springer VS, 2016.
- [14] P. Trautman and A. Krause, "Unfreezing the robot: Navigation in dense, interacting crowds," in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2010, pp. 797–803.
- [15] G. Galati, S. Primates, S. Grammatico, S. Macrì, and A. Rizzo, "Game theoretical trajectory planning enhances social acceptability of robots by humans," *Scientific reports*, vol. 12, no. 1, p. 21976, 2022.
- [16] Helbing and Molnár, "Social force model for pedestrian dynamics," *Physical review. E, Statistical physics, plasmas, fluids, and related interdisciplinary topics*, vol. 51, no. 5, pp. 4282–4286, 1995.
- [17] J. van den Berg, S. J. Guy, M. Lin, and D. Manocha, "Reciprocal n-body collision avoidance," in *Robotics Research*, ser. Springer Tracts in Advanced Robotics, B. Siciliano, O. Khatib, F. Groen, C. Pradalier, R. Siegwart, and G. Hirzinger, Eds. Springer Berlin Heidelberg, 2011.
- [18] 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)*. IEEE, 2017, pp. 285–292.
- [19] 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.
- [20] T. Fan, X. Cheng, J. Pan, D. Manocha, and R. Yang, "Crowdmov: Autonomous mapless navigation in crowded scenarios," 2018.
- [21] L. Kastner, Z. Shen, C. Marx, and J. Lambrecht, "Autonomous navigation in complex environments using memory-aided deep reinforcement learning," in *2021 IEEE/SICE International Symposium on System Integration (SII)*. IEEE, 2021, pp. 170–175.
- [22] B. Brito, M. Everett, J. P. How, and J. Alonso-Mora, "Where to go next: Learning a subgoal recommendation policy for navigation among pedestrians," 25.02.2021.
- [23] M. Everett, Y. F. Chen, and J. P. How, "Motion planning among dynamic, decision-making agents with deep reinforcement learning," in *2019 International Conference on Robotics and Automation (ICRA)*, 2019, pp. 3052–3059.
- [24] C. Chen, Yuejiang Liu, Sven Kreiss, and Alexandre Alahi, *Crowd-Robot Interaction: Crowd-aware Robot Navigation with Attention-based Deep Reinforcement Learning*. Piscataway, NJ: IEEE, 2019.
- [25] M. Everett, Y. F. Chen, and J. P. How, "Collision avoidance in pedestrian-rich environments with deep reinforcement learning," *IEEE Access*, vol. 9, pp. 10357–10377, 2021.
- [26] C. Chen, S. Hu, P. Nikdel, G. Mori, and M. Savva, "Relational graph learning for crowd navigation," 2020.
- [27] E. T. Hall, *The hidden dimension*. New York: Anchor Books, 1996.
- [28] Christoforos I. Mavrogiannis, Valts Blukis, and Ross A. Knepper, "Socially competent navigation planning by deep learning of multi-agent path topologies," 2017.
- [29] S. Rossi, A. Rossi, and K. Dautenhahn, "The secret life of robots: Perspectives and challenges for robot's behaviours during non-interactive tasks," *International Journal of Social Robotics*, 2020.
- [30] S. Guillén-Ruiz, J. P. Bandera, A. Hidalgo-Paniagua, and A. Bandera, "Evolution of socially-aware robot navigation," *Electronics*, 2023.
- [31] P. T. Singamaneni, P. Bachiller-Burgos, L. J. Manso, A. Garrell, A. Sanfeliu, A. Spalanzani, and R. Alami, "A survey on socially aware robot navigation: Taxonomy and future challenges," *The International Journal of Robotics Research*, 2024.
- [32] N. Ah Sen, P. Carreno-Medrano, and D. Kulić, "Human-aware subgoal generation in crowded indoor environments," in *Social Robotics*, ser. Lecture Notes in Computer Science, F. Cavallo, J.-J. Cabibihan, L. Fiorini, A. Sorrentino, H. He, X. Liu, Y. Matsumoto, and S. S. Ge, Eds. Springer Nature Switzerland, 2022, vol. 13817, pp. 50–60.
- [33] Y. Chen, C. Liu, M. Liu, and B. E. Shi, "Robot navigation in crowds by graph convolutional networks with attention learned from human gaze," 23.09.2019.
- [34] Y. Zhou and J. Garcke, "Learning crowd behaviors in navigation with attention-based spatial-temporal graphs," 11.01.2024.
- [35] S. S. Samsani and M. S. Muhammad, "Socially compliant robot navigation in crowded environment by human behavior resemblance using deep reinforcement learning," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 5223–5230, 2021.
- [36] B. Xue, M. Gao, C. Wang, Y. Cheng, and F. Zhou, "Crowd-aware socially compliant robot navigation via deep reinforcement learning," *International Journal of Social Robotics*, 2023.
- [37] V. Narayanan, B. M. Manoghar, R. P. RV, and A. Bera, "Ewaret: Emotion-aware pedestrian intent prediction and adaptive spatial profile fusion for social robot navigation," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023.
- [38] H. Yang, C. Yao, C. Liu, and Q. Chen, "Rmrl: Robot navigation in crowd environments with risk map-based deep reinforcement learning," *IEEE Robotics and Automation Letters*, vol. 8, no. 12, 2023.
- [39] A. B. Richard Sutton, "Reinforcement learning: An introduction: second edition," 2020.
- [40] 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.
- [41] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," 20.07.2017.
- [42] M. Towers, J. K. Terry, A. Kwiatkowski, J. U. Balis, G. de Cola, T. Deleu, M. Goulão, A. Kallinteris, KG, Arjun, M. Krimmel, R. Perez-Vicente, A. Pierré, S. Schulhoff, J. J. Tai, A. T. J. Shen, and O. G. Younis, "Gymnasium," 2023.
- [43] Antonin Raffin, Ashley Hill, Adam Gleave, Anssi Kanervisto, Maximilian Ernestus, and Noah Dormann, "Stable-baselines3: Reliable reinforcement learning implementations," *Journal of Machine Learning Research*, vol. 22, no. 268, pp. 1–8, 2021.