



UvA-DARE (Digital Academic Repository)

Social Navigation with Human Empowerment Driven Deep Reinforcement Learning

van der Heiden, T.; Mirus, F.; van Hoof, H.

DOI

[10.1007/978-3-030-61616-8_32](https://doi.org/10.1007/978-3-030-61616-8_32)

Publication date

2020

Document Version

Author accepted manuscript

Published in

Artificial Neural Networks and Machine Learning – ICANN 2020

License

CC BY

[Link to publication](#)

Citation for published version (APA):

van der Heiden, T., Mirus, F., & van Hoof, H. (2020). Social Navigation with Human Empowerment Driven Deep Reinforcement Learning. In I. Farkaš, P. Masulli, & S. Wermter (Eds.), *Artificial Neural Networks and Machine Learning – ICANN 2020: 29th International Conference on Artificial Neural Networks, Bratislava, Slovakia, September 15–18, 2020 : proceedings* (Vol. II, pp. 395-407). (Lecture Notes in Computer Science; Vol. 12397). Springer. https://doi.org/10.1007/978-3-030-61616-8_32

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (<https://dare.uva.nl>)

Social Navigation with Human Empowerment driven Deep Reinforcement Learning

Tessa van der Heiden¹, Florian Mirus¹, and Herke van Hoof²

¹ BMW Group, Landshuter Str. 26, 85716 Unterschleißheim tessa.heiden@bmw.de
florian.mirus@bmwgroup.de

² University of Amsterdam, Science Park, 1098 XH Amsterdam
h.c.vanhoof@uva.nl

Abstract. Mobile robot navigation has seen extensive research in the last decades. The aspect of collaboration with robots and humans sharing workspaces will become increasingly important in the future. Therefore, the next generation of mobile robots needs to be socially-compliant to be accepted by their human collaborators. However, a formal definition of compliance is not straightforward. On the other hand, empowerment has been used by artificial agents to learn complicated and generalized actions and also has been shown to be a good model for biological behaviors. In this paper, we go beyond the approach of classical Reinforcement Learning (RL) and provide our agent with intrinsic motivation using empowerment. In contrast to self-empowerment, a robot employing our approach strives for the empowerment of people in its environment, so they are not disturbed by the robot’s presence and motion. In our experiments, we show that our approach has a positive influence on humans, as it minimizes its distance to humans and thus decreases human travel time while moving efficiently towards its own goal. An interactive user-study shows that our method is considered more social than other state-of-the-art approaches by the participants.

Keywords: reinforcement learning · empowerment · human-robot interaction

1 Introduction

Recent advances in sensor and control technologies have allowed the development of robots assisting people in domestic, industrial, and traffic environments. One key challenge in such settings, where humans and robots share the same workspace, is that the robot must plan safe, collision-free paths, which have to be socially compliant for the humans to accept robots as collaborators in the long run.

However, moving and thereby interacting with people requires robots to follow specific unwritten social rules, for instance, politely keeping their distance [29], which depends not only on the situation [27] but also the social context and the people involved [23]. According to Kruse *et al.* [17], the three main requirements for a robot to navigate in a socially compliant way are *comfort*,




Fig. 1: Our social compliant robot (SCR) uses occupancy maps centered around each human to compute human empowerment, so it minimally disturbs people to pursue their goals.

naturalness, and *sociability*. Robots acting according to these rules will have a higher chance of acceptance by human users.

Existing methods for social robot navigation either model social conventions between agents explicitly [22], or implicitly learn them through Imitation Learning (IL) [26], or even through Reinforcement Learning (RL) [4]. However, explicitly defining rules or reward functions for social navigation is not straightforward, while generating a sufficiently large body of training examples for imitation learning can be cumbersome or infeasible.

Empowerment [16] allows the agent to generate rewards by itself and offers a useful alternative. It is the channel capacity between actions and future states and maximizes an agent’s influence on its near future. In contrast to self-empowerment, where a self-empowered agent will try to push others away to maximize its future rewards, an agent who strives for others’ empowerment maintains the influence of them on their futures [24].

In this paper, we propose a novel approach to social robot navigation employing a combination of Reinforcement Learning (RL) and human empowerment to provide a robot with intrinsic motivation. An agent employing our approach strives for people’s empowerment to minimize disturbance when pursuing their goals and respect people’s personal space. Our contribution is to use the concept of human empowerment introduced by Salge and Polani [24] as an intrinsic reward function for RL for social navigation.

In an extensive evaluation in a simulation environment, we compare our method with state-of-the-art robotic navigation methods. Inspired by [17], we use two additional metrics, the distance between human and robot and robot’s jerk, assessing social behavior to evaluate our approach. Additionally, we study the robot’s influence on people and vice-versa by introducing two new metrics, the travel time and distance of humans and the robot. Last, we assess our approach in a user-study. Our experiments show that our approach can achieve socially compliant robot navigation by using human empowerment. Finally, our method applies to any multi-agent system that requires a robot to interact with humans in a socially compliant way, since it does not require a cost function.

2 Related work

Many approaches have designed interaction models to enhance social awareness in robot navigation. We discuss these methods first and motivate the practicality of Deep Reinforcement Learning (DRL). We proceed by describing empowerment as a member of a family of intrinsic motivators for RL.

2.1 Social navigation

The goals for social navigation can be divided into three main categories: *comfort*, *naturalness* and *sociability* [17]. Examples of *comfort* are respecting personal space, avoiding erratic behavior, and not interfering with the other’s movement. *Naturalness* is mostly related to how similar a robot’s motion is to human behavior, e.g., smooth and interpretable, while *sociability* is mainly associated with social conventions and etiquettes. Previous works have tried to create navigation frameworks satisfying all of those requirements. Well-engineered methods are the Social Force Model (SFM) [11], Interacting Gaussian Process (IGP) [30], Optimal Reciprocal Collision Avoidance (ORCA) [12] and Reciprocal Velocity Obstacles (RVO) [32]. Although all of these methods yield collision-free paths, they rely on manually engineered models with limited additional social characteristics.

In contrast to such model-based approaches, Deep Learning (DL) models have shown to produce more human-like paths [10]. For instance, Deep Neural Networks (DNNs) allow policies to better comply with humans’ social rules [7]. Early works separate the prediction of the environment and the policy’s planning task into two neural networks [2], which could cause the freezing robot problem since the predicted human motion could occupy all the available future space [30].

Imitation Learning (IL) and Inverse Reinforcement Learning (IRL) obtain policies directly from demonstrations [22], which requires an extensive data set due to the uncertainty of human motion. Alternatively, DRL aims to learn cooperative strategies by interacting with the environment [6]. However, the definition of a suitable cost function that encourages a robot to navigate socially is a challenging task. Even if a cost function might appear evident in some cases (e.g., collision-free and keeping distance to neighbors), it often has to be regularised to achieve smooth and risk-averse behavior.

2.2 Empowerment

Instead of shaping the reward function to achieve the desired behavior, an emerging field within RL focuses on intrinsic motivation [21]. There are many different ways to motivate an agent intrinsically, and one possible technique is called empowerment [15], [25]. Empowerment was applied to teach agents task-independent behavior and training in settings with sparse rewards, such as stabilizing an inverted pendulum, learning a biped to walk [13]. Aubret *et al.* [1] provides a comprehensive survey of empowerment for RL.

Earlier approaches were only applicable in discrete state-action spaces, but recently [20] show efficient implementations for continuous settings. In our work, we will build upon these models.

3 Methodology

Our goal is to teach an agent on how to navigate its target in a socially compliant manner safely. A combination of two rewards can achieve these two objectives. In this section, we describe our agent and the two types of rewards.

We consider the system to be Markovian, where each next state x_{t+1} depends only on the current state x_t and agent’s action u_t and no prior history. A value network model V_ϕ is trained to accurately approximate the optimal value function V^* that implicitly encodes social cooperation between agents and the empowerment of the people, see Equation 1.

$$V^*(x_t) = \sum_{t=0}^T \gamma^t R_t(x_{t+1}, \pi^*(x_t)) \quad (1)$$

$R_t(\cdot)$ is the reward function and π^* is the optimal policy that maximizes the expected return, with discount factor $\gamma \in (0, 1)$.

3.1 Reward for safe navigation

The first task of the agent is to reach its goal while avoiding collisions and keeping a comfortable distance to humans. We consider the number of humans in each episode to be fixed.

Equation 2 defines the environmental reward function $R_{t,e}(x_t, u_t^\pi)$ for this task with the robot’s state denoted as x_t and its action with u_t^π . Similar to other DRL methods for social navigation [5], [6], [4] we award task accomplishments and penalize collisions or uncomfortable distances.

$$R_{t,e}(x_t, u_t^\pi) = \begin{cases} -.25 & \text{if } d_i < 0.01 \\ -0.1 + \frac{d_i}{2} & \text{else if } 0.01 \leq d_i \leq 0.2 \\ 1 & \text{else if } d_g \leq 0.01 \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

Here $d_g = \|\bar{p} - \bar{p}_g\|_2$ is the robot’s distance to the goal during a time interval Δt and $d_i = \|\bar{p} - \bar{p}_i\|_2$ is the robot’s distance to neighbor i . It gets rewarded when its current position \bar{p} reaches the position of the goal \bar{p}_g , but penalized if its position is too close to another one’s position \bar{p}_i .

The robot’s own state, \mathbf{x} , consists of a 2D position vector $\bar{p} = [p_x, p_y]$ and 2D velocity vector $\bar{v} = [v_x, v_y]$. The human states are denoted by $\mathbf{X} = [X_1, X_2, \dots, X_k]$, which is a concatenated vector of states of all k humans participating in the scene. Each entry is similar to the robot’s state, namely, $X_i = [\bar{p}_i, \bar{v}_i]$. The final state of the robot is the concatenation of the state of the humans and robot, $x_t = [\mathbf{X}_t, \mathbf{x}_t]$. Its action is a desired velocity vector, $u_t^\pi = \bar{v}_d$

3.2 Empowerment for social compliance

The robot’s second task is to consider people in its neighborhood and respond to their intentions in a socially compliant manner. Designing a suitable reward function is a challenging task, among other things, due to the stochasticity in people’s behaviors. This is where we use empowerment [15], [25], an information-theoretic formulation of an agent’s influence on its near future.

Human empowerment Empowerment in our case, motivates the robot to approach states in which its neighbors are most empowered. Now the robot aims to maximize the empowerment of another person rather than its own, which Salge and Polani [24] call *human empowerment* in contrast to *robot empowerment*. As a result, the robot will prevent obstructing the human, for example, by getting too close or by interfering with the human’s actions, both of which Kruse *et al.* [17] defined as social skills.

Equation 3 describes the definition of empowerment ε being the maximal mutual information I for a state z_t [15]. It is the channel capacity between action u_t^ω and future state z_{t+1} , maximized over source policy ω . Policy ω is part of the human’s decision making system.

$$\varepsilon(z_t) = \max_{\omega} I(z_{t+1}, u_t^\omega | z_t) = \max_{\omega} H(u_t^\omega | z_t) - H(u_t^\omega | z_t, z_{t+1}) \quad (3)$$

The right part defines the empowerment with entropies $H(\cdot)$. It corresponds to increasing the diversity of decisions, while at the same time limiting those decision that have no effect. Intuitively, the empowerment of the person reflects his or her ability to influence their future.

The human state z_t takes an ego-centric parameterization [5]. Each state is an occupancy grid map centered around the person, denoted with \mathbf{g}_k . It is a 3D tensor with dimensions $c \times r \times 3$, where c and r run over the height and width of the grid. Each entry contains the presence and velocity vector of a neighbor j at that location $\bar{e}_j = [1, v_x, v_y]$. The resulting state of the humans is a concatenated vector denoted by $z_t = [\mathbf{g}_1, \mathbf{g}_2, \dots, \mathbf{g}_k]$ and action are continuous values in \mathbb{R}^2 .

Estimating empowerment with neural networks To compute empowerment, we consider the mutual information defined by the Kullback-Leibler divergence [18] between the joint $p(z_{t+1}, u_t | z_t)$ and the product of the marginal distributions $p(z_{t+1} | z_t)$ and $\omega(u_t | z_t)$:

$$\begin{aligned} I &= D_{KL}(p(z_{t+1}, u_t | z_t) \parallel p(z_{t+1} | z_t) \omega(u_t | z_t)) \\ &= \iint p(z_{t+1}, u_t | z_t) \ln \frac{p(z_{t+1}, u_t | z_t)}{p(z_{t+1} | z_t) \omega(u_t | z_t)} dz' du \end{aligned} \quad (4)$$

The main problem in the formulation in Eq. 4 is the intractability due to the integral of all future states. Since the introduction of empowerment [15], [25] many have designed methods to deal with this. Recent works provide an efficient

method to estimate a lower bound on empowerment \hat{I} , via variational methods [13], [20], [3].

$$\hat{I} = \iint p(z_{t+1}, u_t | z_t) \ln \frac{q(u_t | z_{t+1}, z_t)}{\omega(u_t | z_t)} dz' du \quad (5)$$

Instead of $p(z_{t+1} | z_t)$ a planning distribution $p(u_t | z_t, z_{t+1})$ is used, which is approximated with the variational approximation $q(u_t | z_{t+1}, z_t)$ to obtain a lower bound. \hat{I} can now be maximized over the parameters of the source, $\omega(u_t | z_t)$ and variational $q(u_t | z_{t+1}, z_t)$ networks. $p(u_t | z_t, z_{t+1})$ is a third neural network that computes the future state z_{t+1} from z_t and u_t^ω .

The gradient can be computed as follows, in which the joint parameters of $\omega(\cdot)$ and $q(\cdot)$ is denoted by θ :

$$\frac{\partial}{\partial \theta} \hat{I} = \frac{\partial}{\partial \theta} \mathbb{E}_{p(z_{t+1}, u_t | z_t)} \left[\ln \frac{q(u_t | z_{t+1}, z_t)}{\omega(u_t | z_t)} \right] \quad (6)$$

Using Monte-Carlo integration to estimate the continuous case, we can obtain the following gradient:

$$\frac{\partial}{\partial \theta} \hat{I} \approx \sum_{n=1}^N \frac{\partial}{\partial \theta} [\ln(q(u_t | z_{t+1}, z_t)) - \ln(\omega(u_t | z_t))] \quad (7)$$

We are free to choose any type of distribution and since human movement is not discrete, we model both $q(u_t | z_{t+1}, z_t) = \mathcal{N}(\mu_q, \sigma_q^2)$, $p(z' | u, z) = \mathcal{N}(\mu_p, \sigma_p^2)$ and $\omega(u_t | z_t) = \mathcal{N}(\mu_\omega, \sigma_\omega^2)$ as Gaussian distributions.

3.3 Training procedure

The robot with policy π learns to safely navigate to its goal and achieve human empowered states. This is achieved by training a value network V_ϕ with the reward function combining the mutual information $\hat{I}_t(\cdot)$ and the environmental reward $R_{t,e}(\cdot)$. The hyper-parameter β is used to regulate the trade-off between social compliance and safety:

$$R_t(z_t, x_t, u_t^\omega, u_t^\pi) = (1 - \beta) \cdot I_t(z_{t+1}, u_t^\omega | z_t) + \beta \cdot R_{t,e}(x_t, u_t^\pi) \quad (8)$$

A set of demonstrations from the ORCA policy is used to give the robot a head start. ORCA describes a deterministic control policy [32] and its demonstrations speed up learning, because experiences in which the robot reaches the goal are now part of the memory. Next, the behavior policy π collects samples of experience tuples $e_t = (z_t, x_t, u_t^\pi, r_{t,e}, z_{t+1}, x_{t+1})$ until a final state is reached. Random actions are selected with probability ϵ . Once these are collected, our hypothetical human policy ω together with q and p are used to estimate \hat{I}_t . Finally, the networks are trained with a random mini-batch obtained from the memory (e^b). The value network V_ϕ is optimized by the Temporal Difference

Method (TDM) [28] with standard experience replay and fixed target network techniques [5], [19], [6].

$$\begin{aligned} y_t &= r_{e,t}^{(b)} + \hat{I}_t^{(b)} + \gamma^t \hat{V}_{\hat{\phi}}(z_{t+1}^{(b)}, x_{t+1}^{(b)}) \\ \phi &\leftarrow \phi + \lambda \nabla_{\phi} (y_t - V_{\phi}(z_t, x_t))^2 \end{aligned} \quad (9)$$

$\hat{V}_{\hat{\phi}}$ denotes the target network. The networks ω_{θ} and q_{θ} updated through gradient ascent and p_{ψ} via gradient descent:

$$\begin{aligned} \theta &\leftarrow \theta + \lambda \nabla_{\theta} \hat{I}_t \\ \psi &\leftarrow \psi + \lambda \nabla_{\psi} (z_{t+1}^{(b)} - p(u_t, z_t))^2 \end{aligned} \quad (10)$$

The behavior policy π uses the joined state x_t to navigate collision-free to its goal. p , q and ω take the occupancy grids centered around each human z_t as states for the computation of \hat{I}_t .

4 Experiments

We conduct three experiments to evaluate our proposed model. The first experiment compares our model against four existing state-of-the-art methods. The second experiment assess the social competences, based on the metrics defined by Kruse *et al.* [17]. The final experiment consists human subjects that evaluate the models in an interactive simulator.

4.1 Implementation details

The simulator used in this work is from [6]. It starts and terminates an episode with five humans and the robot. The human’s decisions are simulated by van den Berg *et al.* [31], which uses the ORCA policy [32] to calculate their actions. ORCA uses the optimal reciprocal assumption to avoid other agents. In 4.4, another simulator is used to control the position of one human manually and terminates once the human reaches its goal. The code and videos can be found online³.

We implemented the networks in PyTorch and trained them with a batch size of 100 for 10 000 episodes. For the value network, the learning rate is $\lambda_v = 0.001$ and the discount factor γ is 0.9. The exploration rate of the ϵ decays linearly from 0.5 to 0.1 in the first 5000 episodes and stays 0.1 for the remaining 5000 episodes. These values are the same as Chen *et al.* [6].

The parameter β is 0.25, because that gave the highest discomfort distance rate and success rate. The learning rates for the other networks are similar. The value network is trained with stochastic gradient descent, identical to Chen *et al.* [6]. The planning, source, and transition networks are trained with Adam [14], similar to [13].

³ <https://github.com/tessavdheiden/SCR>

4.2 State-of-the-art navigation benchmark

Table 1 reports the rates of success, collision, the robot navigation time, discomfort distance, and the average discounted cumulative reward averaged over 500 episodes. **Success** is the rate of robot reaching its goal within a time limit of 20s. **Collision** is the rate of the robot colliding with humans. **Disc. dist.** is the rate in which the distance between the robot and a human was smaller than 0.1. We compare our robot with four existing state-of-the-art methods, ORCA [31], Collision Avoidance with Deep Reinforcement Learning (CADRL) [5], Long Short Term Memory - Reinforcement Learning (LSTM-RL) [8] and Socially Attentive Reinforcement Learning (SARL) [6]. As can be seen, our Socially Compliant Robot (SCR) and SARL both outperform other baselines on the standard metrics. Next, we look more thoroughly into the performance of SCR and SARL.

Table 1: Both SCR and SARL outperform the other baselines, which can be seen by the best values (grey). ORCA does not have any collisions, because this is the central idea behind the method (*). The numbers are computed for 500 different test scenarios.

	Success	Collision	Robot	Disc.	Reward
	rate %	rate %	time	dist. %	
ORCA	0.99	.000*	12.3	0.00*	.284
CADRL	0.94	.035	10.8	0.10	.291
LSTM-RL	0.98	.022	11.3	0.05	.299
SARL	0.99	.002	10.6	0.03	.334
SCR (ours)	0.99	.001	10.9	0.03	.331

4.3 Influence of robot on humans and vice-versa

Table 2 shows travel times and distances of both humans and the robot. **Robot time** and **Human time** are the average navigation times of the agents to reach their goals. Keeping both low indicates that the robot does not disturb humans and moves quickly to its target. The path length of the robot, **Trav. distance**, is calculated to make sure that it moves without unnecessary detours. The simulator allows making the robot invisible to the humans, which is called the Invisible baseline (Invisible Baseline (IB)). This baseline serves as a testbed for validating the other policies’ abilities in reasoning about the interactions with humans.

IB has the highest travel time and distance because it does not influence humans. On the contrary, SARL has a low travel distance and time, but human travel times are highest. SARL has learned that humans avoid it. The travel times of SCR and that of the humans are nearly the same. These numbers suggest that our method has learned to minimally disturb other people while moving to its own goal efficiently due to human empowerment.

Next, we examine how we can evaluate social compliance further. Fong *et al.* [9] and Kruse *et al.* [17] state that people judge robots negatively if the separation distance between them is low and move non-smoothly. **Sep. distance** is the distance between a human and the robot and **Jerk** is the jerk of the robot ($\frac{m}{s^3}$).

Table 2: SCR moves efficiently to its target and doesn’t disturb people as their travel times are low. SARL has learned that people avoid it, so the humans’ travel times are higher than its own. The IB takes a considerable detour because people do not see it. Moreover, SCR has the lowest jerk ($\frac{m}{s^3}$), since it avoids being close to a person and non-smooth behavior as this would lower the empowerment of its neighbors.

	Robot	Trav.	Human	Sep.	Jerk
	time	distance	time	distance	
IB	11.5	10.7	9.1	.85	.73
SARL	10.6	9.2	10.7	.41	.77
SCR (ours)	10.9	9.3	9.1	.43	.51

The reported numbers are in the last two columns of Table 2. Even though on average SCR and SARL are close to humans, they do not exceed a minimum distance of .1m, see collision rate in Table 1. The IB and SARL move with a high jerk because their reward functions do not incorporate it. On the other hand, SCR has the lowest jerk because it avoids erratic behavior as that would lower humans’ empowerment.

4.4 Human evaluation

A successful navigation strategy can best be tested with real persons. To that end, 30 persons interacted with the robot and controlled the position of one human in the simulator ³. The robot and humans start at 90° from each other and need to cross the center to reach their goals. The simulator terminates once the human reaches his/her goal. After that, subjects are asked to rate the social performance of the robot with a score from 1 to 10 (similar to the study performed by [26]). A one-way repeated-measure analysis of variance (ANOVA) was conducted to determine significance, and a significant effect was obtained ($f(2, 27) = 13.485, p < .01$).

As can be seen in Fig. 2, both SARL and the IB have similar medians, but the samples of the baseline deviate more from the median. The IB started to move away from the human, even before the human started to move. SARL, on the other hand, moves directly to its own goal. SCR has the highest score, which shows the potential of our method.

4.5 Qualitative results

Figure 3 shows SARL and SCR navigating through a crowd of five people. The left figures shows SARL (a, b) and right SCR (c, d) at two different time steps. The trajectories indicate that SARL goes directly to its goal, while SCR waits at $t = 6$ (c). Moreover, at $t = 9.2$, SARL has reached its goal, but only two out of five humans reach theirs (b, purple and light blue stars). In contrast, SCR reaches its goal at $t=10.5$, but all people reached their final destinations (d). SARL overtakes two people (a, red and green) and alters the path of another (a, blue). On the contrary, SCR lets them pass (c, red, green and blue). SARL uses occupancy maps to model the pairwise interaction between humans [6], so it

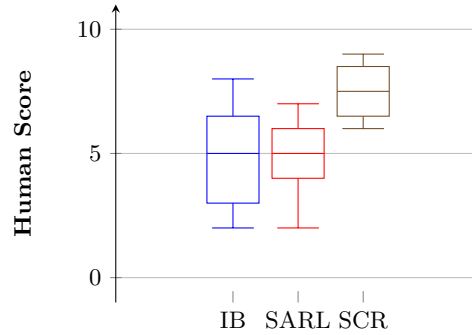


Fig. 2: The box plots summarize human evaluation scores on three methods, the IB, SARL and our method (SARL). Both the IB and SARL have almost the same median, but for the IB the humans’ scores deviate more from the median. Our method obtains the highest score.

cannot incorporate the robot’s influence on each human. On the contrary, SCR uses empowerment maps for each human with high values in states in which it does not block anyone.




Fig. 3: SARL (a, b) and SCR (c, e) in a scene with 5 humans. The humans’ destinations are the opposite of the x, y-axis’ origin from their initial locations. SARL reaches its destination quickly, but only two out of five humans reach it (b, 2 stars). Two persons (red, blue) need to adjust their path to avoid the robot (orange). SCR waits at $t=6$ (c) and all humans reach their destination (d, 5 stars).

5 Conclusion and future work

This paper proposed a reinforcement learning method for social navigation with intrinsic motivation using the empowerment of surrounding people. Our approach avoids the hard-coded reward signals and allows people nearby not to be disturbed by the robot. Our experiments show that our policy outperforms other methods on social characteristics. The influence of the robot’s motion is difficult to evaluate by people in simulation. Thus, we also compared the methods in an interactive simulator and obtained positive results. For future work, we would like to extend the model to deal with a variable amount of humans and with different policies. It would also be interesting to extend the method to incorporate the effect that (non-moving) objects have on humans.

Bibliography

- [1] A. Aubret, L. Matignon, and S. Hassas. A survey on intrinsic motivation in reinforcement learning. *arXiv preprint arXiv:1908.06976*, 2019.
- [2] S Bansal, V Tolani, S Gupta, J Malik, and C Tomlin. Combining optimal control and learning for visual navigation in novel environments. *arXiv preprint arXiv:1903.02531*, 2019.
- [3] Y. Burda, R. Grosse, and R. Salakhutdinov. Importance weighted autoencoders. *arXiv preprint arXiv:1509.00519*, 2015.
- [4] Y. Chen, M. Everett, M. Liu, and J. P. How. Socially aware motion planning with deep reinforcement learning. In *2017 IEEE/RSJ IROS*, pages 1343–1350. IEEE, 2017.
- [5] Y. Chen, M. Liu, M. Everett, and J. P. How. Decentralized non-communicating multiagent collision avoidance with deep reinforcement learning. In *2017 IEEE ICRA*, pages 285–292. IEEE, 2017.
- [6] C. Chen, Y. Liu, S. Kreiss, and A. Alahi. Crowd-robot interaction: Crowd-aware robot navigation with attention-based deep reinforcement learning. In *2019 ICRA*, pages 6015–6022. IEEE, 2019.
- [7] E. Cross, R. Hortensius, and A. Wykowska. From social brains to social robots: applying neurocognitive insights to human-robot interaction. *Philosophical Transactions of the Royal Society of London. Series B, Biological sciences*, 374, 04 2019.
- [8] M. Everett, Y. Chen, and J.P. How. Motion planning among dynamic, decision-making agents with deep reinforcement learning. In *2018 IEEE/RSJ IROS*, pages 3052–3059. IEEE, 2018.
- [9] T. Fong, I. Nourbakhsh, and K. Dautenhahn. A survey of socially interactive robots. *Robotics and autonomous systems*, 42(3-4):143–166, 2003.
- [10] T. Gu and J. Dolan. Toward human-like motion planning in urban environments. In *2014 IEEE Intelligent Vehicles Symposium Proceedings*, pages 350–355. IEEE, 2014.
- [11] D. Helbing and P. Molnar. Social force model for pedestrian dynamics. *Physical review E*, 51(5):4282, 1995.
- [12] I Karamouzas, P Heil, P. Van Beek, and M.H. Overmars. A predictive collision avoidance model for pedestrian simulation. In *International workshop on motion in games*, pages 41–52. Springer, 2009.
- [13] M. Karl, M. Soelch, P. Becker-Ehmck, D. Benbouzid, P. van der Smagt, and J. Bayer. Unsupervised real-time control through variational empowerment. *arXiv preprint arXiv:1710.05101*, 2017.
- [14] D.P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [15] A.S. Klyubin, D. Polani, and C.L. Nehaniv. Empowerment: A universal agent-centric measure of control. In *2005 IEEE Congress on Evolutionary Computation*, volume 1, pages 128–135. IEEE, 2005.

- [16] A.S. Klyubin, Daniel Polani, and Chrystopher Nehaniv. Empowerment: A universal agent-centric measure of control. volume 1, pages 128 – 135 Vol.1, 10 2005.
- [17] T. Kruse, A.K. Pandey, R. Alami, and A. Kirsch. Human-aware robot navigation: A survey. *Robotics and Autonomous Systems*, 61(12):1726–1743, 2013.
- [18] Solomon Kullback and Richard A Leibler. On information and sufficiency. *The annals of mathematical statistics*, 22(1):79–86, 1951.
- [19] 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*, 518(7540):529, 2015.
- [20] S. Mohamed and D.J. Rezende. Variational information maximisation for intrinsically motivated reinforcement learning. In *NeurIPS*, pages 2125–2133, 2015.
- [21] P Oudeyer, F. Kaplan, and V.V. Hafner. Intrinsic motivation systems for autonomous mental development. *IEEE Transactions on Evolutionary Computation*, 11(2):265–286, 2007.
- [22] 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 *2016 IEEE/RSJ IROS*, pages 2096–2101. IEEE, 2016.
- [23] A. Robicquet, A. Sadeghian, A. Alahi, and S. Savarese. Learning social etiquette: Human trajectory understanding in crowded scenes. In *European conference on computer vision*, pages 549–565. Springer, 2016.
- [24] C. Salge and D. Polani. Empowerment as replacement for the three laws of robotics. *Frontiers in Robotics and AI*, 4:25, 2017.
- [25] C. Salge, C. Glackin, and D. Polani. Empowerment—an introduction. In *Guided Self-Organization: Inception*, pages 67–114. Springer, 2014.
- [26] Kyriacos Shiarlis, Joao Messias, and Shimon Whiteson. Acquiring social interaction behaviours for telepresence robots via deep learning from demonstration. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 37–42. IEEE, 2017.
- [27] A. Sieben, J. Schumann, and A. Seyfried. Collective phenomena in crowds—where pedestrian dynamics need social psychology. *PLoS one*, 12(6):e0177328, 2017.
- [28] R.S. Sutton, A.G. Barto, et al. *Introduction to reinforcement learning*, volume 2. MIT press Cambridge, 1998.
- [29] A. Templeton, J. Drury, and A. Philippides. Walking together: behavioural signatures of psychological crowds. *Royal Society open science*, 5(7):180172, 2018.
- [30] P. Trautman and A. Krause. Unfreezing the robot: Navigation in dense, interacting crowds. In *2010 IEEE/RSJ IROS*, pages 797–803. IEEE, 2010.
- [31] J. van den Berg, S.J. Guy, J. Snape, M.C. Lin, and D. Manocha. Rvo2 library: Reciprocal collision avoidance for real-time multi-agent simulation.
- [32] J. Van den Berg, M. Lin, and D. Manocha. Reciprocal velocity obstacles for real-time multi-agent navigation. In *2008 IEEE International Conference on Robotics and Automation*, pages 1928–1935. IEEE, 2008.