Context-appropriate Social Navigation in Various Density Construction Environment using Reinforcement Learning

YeSeul Kim [a], Bogyeong Lee[b], Robin Murphy[c], and Changbum R. Ahn[d]

[a] Department of Multidisciplinary Engineering, Texas A&M University, College Station, TX 77843
[b] Department of Architectural Engineering, Dankook University, Korea
[c] Department of Computer Science and Engineering, Texas A&M University, College Station, TX 77843
[d] Department of Architecture and Architectural Engineering, Seoul National University, Seoul, South Korea 08826
E-mail: leslieyskim@tamu.edu, bglee_@dankook.ac.kr, robin.r.murphy@tamu.edu, cbahn@snu.ac.kr

Abstract –

Construction environments are often densely populated with multiple resources (e.g., workers, equipment, and materials). As an increasing number of mobile robots are expected to coexist and interact with humans at close proximity, it is necessary that these robots are capable of not only avoiding collisions with people but also not disturbing human work and deteriorating human comfort. Failing to maintain a proper social space can lead to fatal accidents and inefficiency. To accommodate this need, this study aims to develop a social navigation model that enables robots to navigate in a contextually compliant manner. We created a simulation environment where robot agents can learn socially and contextually aware policies using reinforcement learning. The results showed that the agent was able to secure the respective minimum separation distance for different types of workers while achieving similar overall performance in contrast to baseline models which often violated the work-related proxemic considerations. This finding will contribute to building future construction mobile robots with social intelligence which are capable of understanding the context of the workplace and adapting to appropriate behaviors accordingly.

Keywords –

Construction mobile robots; Social Navigation, Human-Robot Interaction, Reinforcement Learning

1 Introduction

Construction environments are often densely populated with a variety of resources such as workers, equipment, and materials, and they are usually operated in close proximity. As a result of the environment's congested and dynamic nature, the construction industry suffers from contact collisions between workers and equipment [1]. To prevent high traffic accidents, it has been essential to maintain a certain distance from one another.

As a growing number of mobile robotics engaging in non-permanent construction tasks such as reality capture [2,3], safety surveillance [4], and environment monitor (e.g., illuminance measurement) [5] are developed and expected to be deployed in near future, the introduction of mobile robots will bring new safety risks to workers who are in vicinity to the robot trajectory. During operations, robots will encounter a number of workers in different situations, and their movement and the resulting interaction with people can create discomfort and safety threats, causing rejection [6].

Thus, to deploy autonomous mobile robots in construction sites, safe and efficient navigation that produces socially acceptable robot behaviors in the context of construction sites is a vital precondition. This means that these robots are capable of moving through crowds of people while preserving a minimum distance from the co-existing people. This brings the notion of social navigation, which accounts for social conventions such as comfort, naturalness, and sociability, in addition to traditional navigation objectives such as obstacle avoidance and task completion. Due to its importance, robot mobility in a socially compliant manner has been an active area of research in various domains [7,8].

However, one limitation of the existing socially aware navigation models is the lack of understanding of individuals' different proxemic requirement based on their work context. Previous research studies treated each individual equally in their models, assuming no difference in proxemic requirements or preferences among different groups of pedestrians [9–11]. In construction environment, nevertheless, individual pedestrian workers possess unique proxemic considerations based on their work-related contexts. Each worker requires different personal spaces depending on his operational status and space constraints. For example, if a worker is carrying heavy materials where various safety hazards are presented, he needs a larger separation distance from others to feel safe and uninterrupted. To this end, robots should exhibit appropriate proxemic behaviors with respect to these different needs of workers in order to be seamlessly integrated into the construction environment.

Our study aims to develop a context-appropriate social navigation (CASN) algorithm sensitive to different types of pedestrian workers' proxemic needs in construction sites. We used different reinforcement learning (RL) algorithms on the CASN model and evaluated them against baseline models. The preliminary results demonstrated that the proposed model can navigate in crowds with appropriate social etiquette comparable to state-of-the-art methods with some limitations. This finding will contribute to building future construction mobile robots with social intelligence capable of understanding the workplace context and taking socially and contextually appropriate behaviors accordingly.

2 Background

2.1 Social Spaces and Proximity Considerations in Construction

For navigation tasks in a human-populated environment, mobile robots' behaviors must be acceptable by the humans who share the same space with robots. One paramount aspect of such spatial behaviors is that it does not intrude on people's social space, including personal space and activity space as humans perceive [12].

The notion of personal space is first introduced and delineated in the proxemic theory by Hall [13]. Proxemics is the study of spatial distances that individuals maintain in various social and interpersonal situations, and it is used to define interaction strategies. According to Hall, people can perceive and manage their personal space from others and respect others' space in a similar manner. Similar to personal space, actions performed by humans constitute activity space. Other people maintain this space to avoid disturbing the activity.

These social spaces, such as personal or activity spaces, often depend on environmental or cultural factors. Thus, the preferred social space distance is subject to be changed based on the context [14]. For example, if the potential threat is high, personal space distance tends to get larger. Understanding and respecting this personal space is vital in terms of safety assurance. When personal space interferes, people feel discomfort, and it increases physical safety risks.

In the construction domain, which are often characterized as unstructured and cluttered, respecting proxemics, or the minimum distance from other workers is even more critical. Various construction activities are concurrently operated in proximity, and both workers and equipment constantly change their positions in the construction job sites. Consequently, there are many direct and indirect safety risks of shared spaces with unmanned vehicles in proximity [15–17]. Maintaining the relevant separation distance from workers, therefore, is critical for mobile robots to assure both physical and psychological safety.

For example, when autonomous robots are integrated into the construction sites, any robot behavior that impinges on workers' personal space will be perceived as not only unsafe but also unpleasant. In particular, if activity spaces are threatened while they are performing construction activities, they will lose comfort and sense of safety. Some may even feel disrupted and demotivated because the impact of autonomous navigation behaviors on workers varies depending on the context. For example, workers will perceive smaller personal space for those who are just walking compared to those walking during task operations (e.g., transporting materials from one place to another). Similarly, workers will perceive a larger personal space for those in a safety-critical situation than those in less dangerous situations.

Thus, to avoid interrupting the activities and decreasing productivity in construction, it is critical for robots to understand the status of workers and maintain a corresponding separation distance from the workers. In other words, while the primary personal space needs to be respected in the construction domain, it is also essential to preserve context-specific activity spaces in certain work-related situations.

To this end, construction mobile robots should understand these contextual proxemic considerations and constraints safely and effectively navigate the construction site where robots encounter various workers and work situations. The detailed proxemic considerations are further discussed in Section 3.1.

2.2 Socially-aware Navigation Algorithms

Socially-aware navigation refers to the strategy in robot motion planning which accounts for social norms and conventions regarding space management [8]. It is based on the belief that humans interact with robots similarly to human-human interaction; thus, robots can learn from the social norms observed by humans to acquire more acceptable social behaviors and interactions with humans [18]. Socially compliant robots have the ability to perceive and understand crowd behavior to plan their future trajectory to reach their target destination while maintaining an appropriate distance from other pedestrians [7].

There are two approaches in the socially aware navigation literature. A first approach is a model-based approach, which describes the spatial behaviors in terms of a set of rules (e.g., distance, velocity, acceleration, direction) [19]. Using hand-crafted functions to ensure collision avoidance, it performed well in an austere environment; however, it did not generalize well and lacked the capacity to adapt to various situations and lead to oscillatory paths.

To address these limitations, recent works have focused on learning-based approaches such as deep reinforcement learning [11] and imitation learning [20]. In the Inverse Reinforcement Learning (IRL), or imitation learning framework, agents can learn the underlying reward function from human demonstrations. Deep reinforcement learning (DRL) methods such as Deep Q-Network (DQN), Double DQN (DDQN), Trust Region Policy Optimization (TRPO), and Proximal Policy Optimization (PPO), can learn how to behave by interacting with an environment through a sequence of observation, actions, and rewards. Recent work in robot navigation research, DQN, and TPPO showed state-of-the-art performances [21]. To this end, this study leverages the state-of-the-art RL algorithms, namely DQN and TRPO, to develop CASN social navigation model that accounts for various contextual proxemic needs.

3 Methodology

3.1 Context-appropriate Social Navigation

The social space is highly relevant to the context. Construction workers keep a shared convention of construction workspace in construction sites when they interact/work with other workers, and the relevant social space is determined by the activity or task a worker is taking. Thus, the robot's spatial movement should reflect the construction work context to select the appropriate proxemic requirement for different pedestrian workers. This study defined three levels of proxemic considerations and assigned low, medium, and high proxemic requirements for different worker types, as illustrated in Figure 2.

The low proxemic requirement includes normal pedestrian workers who are not conducting construction activities; thus, they are the most flexible in their spatial movement. The medium proxemic requirement group includes load-carrying pedestrian workers. Because these workers have some physical constraints due to their material, they require a larger space than the regular group. Lastly, the high proxemic requirement groups are those who are actively conducting construction activities. Intruding their activity spaces will lead to inefficiency. Thus, priority was placed on this group with the most significant space.

Based on the principles of proxemic theory, we created the corresponding work zone as concentric boundaries and assigned three levels of work zone radius based on the pedestrian worker type. The small work zone indicates that there are less safety risks and activity disturbance compared to the large work zone. We define the work zone distances specific to the work status as $d_{w,s}$, $d_{w,m}$, $d_{w,l}$ for normal pedestrian, load-carrying pedestrian, and operational activity groups, respectively.

Table 1 Work Zone

Pedestrian Worker Type	Work Zone Size	
Normal Pedestrian	Small	0.2
Load-carrying Pedestrian	Medium	0.3
In-Activity	Large	0.4

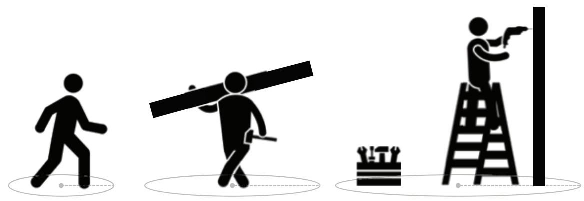

(a) normal-walking, (b) material-carrying walking, and (c) construction activity

Figure 2. Examples of Different Work Zone based on the Working Status

Acknowledging the need for different workers, this study aims to incorporate these different proxemic considerations into social navigation models. We leveraged DQN and TRPO as the base algorithms of the proposed model. Within the RL framework, we formulated this problem as a sequential problem that the robot agent interacts with the environment, observes the states of other humans, and makes a sequence of decisions to maximize the expected return.

DQN is a value-based DRL method that solves the problem by approximating the optimal value function. The optimal policy is to maximize the expected return. It leverages experience replay which stores past experiences and randomly use a subset of them to update the Q-network.

$$V^*(s) = \sum_{t'=t}^{T} \gamma^{t' * v_{pref}} R_{t'}(s_t, \pi^*(s_t)) \tag{1}$$

TPPO, on the other hand, is a policy-based DRL method which becomes very powerful in recent DRL research [22]. Formulating as a hard constraint problem, it aims to maximize the objective function, J(θ), subject to trust-region constraint, which enforces the distance between old and new policies measured by KL-divergence to be small enough with guaranteed monotonic improvement.

$$J(\theta) = E_{s \sim p_{\pi\theta old}, a \sim \pi_{\theta old}} \left[\frac{\pi_\theta(a_t|s_t)}{\pi_{\theta old}(a_t|s_t)} \hat{A}_{\theta old}(s|a) \right] \quad (2)$$

In our model, the radius of the robot is r, and the radius of the human be r_h. The center to center distance between robot and human is defined as d_c. Then, the minimum separation distance between robot and human, d_{min}, is given by

$$d_{min} = d_c - r - r_h \quad (3)$$

To provide feedback to the robot to learn the desired navigation behaviors, the reward function, R_t at timestep t, at state s, and after action a are defined in such a way that the overall navigation goal was fulfilled by achieving a balanced outcome for both success rate and securing the corresponding minimum distance. It implemented different reward functions for different worker groups with corresponding values of work zone spaces based on the operational status. This was designed to be enforced greatly when violating the work zones of worker groups compared to those in regular conditions.

We tried different reward factors for the new reward function, and each reward function calculation is done in the following sequence.

$$R_t(s,a) = \begin{cases} 1, & if\ success \\ (d_{min} - d_{w,s})f, & if\ d_{min} < d_{w,s} \\ (d_{min} - d_{w,m})f, & if\ d_{min} < d_{w,m} \\ (d_{min} - d_{w,l})f, & if\ d_{min} < d_{w,l} \\ 0, & otherwise \end{cases} \quad (4)$$

In addition to negative rewards for collision, we added another set of penalties for invading the assigned personal space for each group. The penalty factor f, is defined for intruding the work zone. This ensures better integrity of the algorithm in the sense that those who are working and situated in safety-critical environments are given higher priority over regular workers.

3.2 Simulation Environment

Because training robots in actual physical environments can be expensive and dangerous, much RL research is done in simulation environments and transfer the learning outcome to the physical environment. In our study, all experiments were carried out using a 2D navigation simulation environment, RVO2 [23]. RVO2 simulates human movements by Optimal Reciprocal Collision Avoidance (ORCA) policy [24]. ORCA uses the optimal reciprocal assumption to generate a path in a shared space while avoiding other agents and preventing collisions with the local observation of the environment, as illustrated in Figure 1.

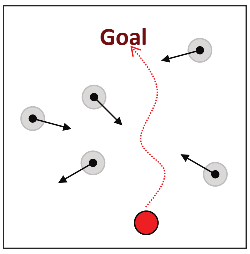

Figure 1. Simulation Environment w/ Reciprocal Collision Avoidance

Our simulation environment is composed of n human and a robot, and the robot task is to navigate toward a given goal in a crowd setting of n pedestrian workers. We assumed that humans could avoid collision with other humans based on ORCA. However, we assumed that humans do not change their paths based on the robot.

The robot, on the other hand, observes the states of the environment, {distance to the goal position}, and adapts its actions to avoid collisions. The robot action space is given {direction, velocity}. It was assumed that the robot knows the goal position, and it can identify the category that a human belongs to in real-time with the vision data it obtains through navigation. The reward function is defined by {Reaching Goal Reward, Collision Reward}.

When the minimum separation distance between robot and human is smaller than the radius value, it would be considered a collision.

The episode will terminate when it reaches the goal in the duration of 25 secs or fails the task by colliding with people or by timeout.

3.3 Experiment

We incorporated the proposed CASN model into DQN and TRPO, and compared them with baseline models. We implemented the ORCA and Socially Attentive Reinforcement Learning (SARL) algorithm as the baseline models. SARL [25] also adopts DQN with an attentive pooling mechanism to learn the importance of neighboring humans with respect to their future trajectories. As it considers various factors like speed and directions of humans as opposed to the previous work assuming the closest neighbors had the most crucial effect on the robot, it improved the performance (e.g., time efficiency) compared to other state-of-the-art benchmark models. These models are used to delineate how the robot would behave without the knowledge of different proxemic considerations.

In addition, we experimented with these algorithms in various density environments. Because robots would encounter a varying density of dynamic obstacles and crowds while moving through the construction sites (e.g., a narrow corridor and/or open space), it is imperative to explore how the density of the crowd would influence the performance of our model. Thus, we compared the results of our models in the same environment with different densities: robot navigating among five people versus ten people to understand how each algorithm would perform differently in various density environments.

For the experiments, we implemented the algorithms in PyTorch. Adopting from the framework of [26] in each set of training experiment, 10,000 training episodes were run with a batch size of 100. We used the learning rate of 0.001 and the discount factor of 0.9. We adopted the ε-greedy policy, where the exploration rate decays linearly from 0.5 to 0.1 in the first 5000 episodes and stays 0.1 for the remaining 5000 episodes.

3.4 Evaluation

As it is challenging to quantify the social conformity of robot behaviors, we used the following five evaluation metrics: success rate, collision rate, navigation duration time, violation of social space, and the minimum distance between human and robot while violation to evaluate social behaviors of the navigation algorithms. Firstly, the success rate is defined as the rate of the robot reaching the target goal within a time limit of 25 seconds, and it is considered as a failure when the robot collides with humans or the total navigation time is over 25 seconds. The collision rate is defined as the rate of robots colliding with humans, which means that the distance between the robot and a human is zero. Violation of social space is defined as the rate at which the minimum distance between the robot and a human is less than the designated work zone size. The minimum distance during violation is defined as the distance between the robot and a human when the robot intruded on the work zone of different workers. We assumed that a larger minimum distance during violation is more socially acceptable.

4 Results

This section compared our context-appropriate social navigation model against the baseline model and showed significant improvements in terms of social conformity. We also describe the performances of proposed models – DQN and TRPO in different density scenarios: high and low density of pedestrian workers.

4.1 Low Density

Our preliminary results showed that the context-appropriate navigation algorithms successfully learned the social norms relative to different workers' proxemic requirements in the low-density simulation environment. Figure 1 show the cumulative discounted rewards of the robot with respect to the number of episodes.

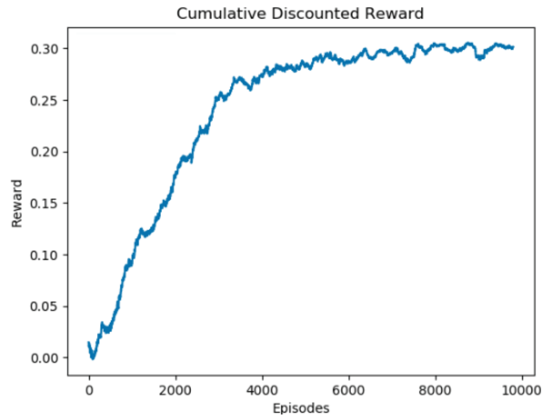

Figure 2. Cumulative discounted rewards for 5 human density

The summarized result is described in Table 2. Except for the ORCA baseline model, most algorithms successfully learned to navigate to the target position. The DQN baseline model showed a relatively similar total navigation duration time. However, in terms of social norm compliance, it suffered from a higher rate of personal space violations and the shorter minimum separation distance during the violation. On the other hand, our CASN models could maintain proper distance for different worker groups, ensuring their physical and psychological safety. It also showed increased minimum distance in case of violation.

In comparing different RL algorithms, the policy gradient algorithm TRPO did not significantly outperform the SARL algorithm in the low-density setting. We observed that all algorithms that we experimented with did not have significant differences in outcomes regarding the success rate, the total navigation time, and frequency of the violation.

Table 2 Comparison of Context-Appropriate Social Navigation model in low density environment

Model	Avg. Duration (s)	Success Rate (Suc no / Total no)	Collision Rate (Col no / Total no)	Frequency of Violation (Violation no. / Total no.)			Avg. Distance during Violation		
				Small	Medium	Large	Small	Medium	Large
ORCA	10.86	0.43	0.57	0.09	0.17	0.10	0.08	0.12	0.17
DQN	10.67	1	0	0.01	0.02	0.03	0.15	0.24	0.34
CASN-DQN	10.55	1	0	0.01	0.01	0.01	0.17	0.25	0.36
CASN-TRPO	10.75	1	0	0.01	0.01	0.01	0.17	0.27	0.38

4.2 High-Density

Similar to those in the low-density scenarios, the preliminary results in the high-density simulation environment also showed that the context-appropriate navigation algorithms successfully learned the social norms relative to different workers' proxemic requirements. However, the overall performance of the CASN models in the high-density environment slightly decreased in terms of total navigation duration, frequency of violation, and the average separation distance during the violation. They tend to converge slowly compared to those in the low-density environment, as illustrated in Figure 2. This tendency of decreased performance was most apparent in the baseline models. They suffered from a higher number of failures and violations, demonstrating the lack of the capacity to handle such high-density situations.

In comparing different RL algorithms, the policy gradient algorithm TRPO outperformed the SARL algorithm in the crowded environment setting, as summarized in Table 3. It reduced the total navigation duration. However, these models were not able to maintain the proper distance from the group with a large personal space requirement.

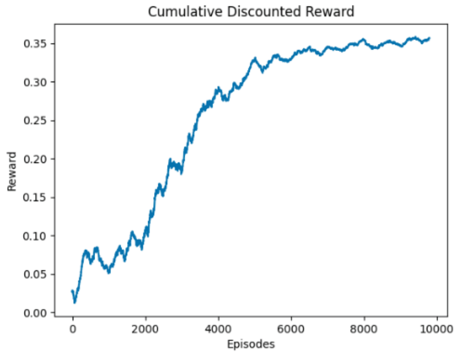

Figure 2. Cumulative discounted rewards for 10 human density

Table 3 Comparison of Context-Appropriate Social Navigation model in high density environment

Model	Avg. Duration (s)	Success Rate (Suc no / Total no)	Collision Rate (Col no / Total no)	Frequency of Violation (Violation no. / Total no.)			Avg. Distance during Violation		
				Small	Medium	Large	Small	Medium	Large
ORCA	12.49	0.21	0.79	0.10	0.14	0.26	0.07	0.13	0.17
DQN	13.25	0.67	0.24	0.07	0.11	0.20	0.13	0.24	0.33
CASN-DQN	12.89	1	0	0.01	0.01	0.02	0.14	0.25	0.35
CASN-TRPO	12.51	1	0	0.01	0.01	0.02	0.15	0.25	0.36

5 Discussion

Our results demonstrated the potentials of the context-appropriate navigation model in making the robot aware of and responsive to varying types of pedestrian workers. Regardless of the RL algorithms, both value-based and policy-based methods showed acceptable outcomes. On the contrary, the baseline models without the proxemic considerations were not able to conform to the varied proxemic requirements, especially for the activity group with large personal space requirement. This suggests that social navigation algorithms without taking consideration of different proximity requirements would cause discomfort particularly to those who are in the most significant need of social norm conformity and in the most serious safety risks.

Although context-appropriate navigation model accounting for proxemics performed well in low to moderate density environments, the performance tends to degrade as the density of the human increase. It showed underperformance in highly dense environments, causing the violation of social spaces and reduced minimum separation distance. This finding is similar to the prevalent limitation of the existing socially-aware navigation models [11]. Trautman et al. [27] showed that it is difficult to avoid freezing robot problem, which refers to a situation where robot halts or oscillates and results in either in a collision or no progress toward the goal, beyond a certain density environment. This finding implies that it will be even more challenging to learn and adhere to the different proxemic needs of workers in situations where workers and moving equipment are densely populated although the consequences of violation in highly dense areas can be more severe and dangerous. These results imply that the current proxemic-based reactive planner may not be effective or safe in the high-density situations. Instead, an interactive and cooperative planner which understands and incorporates context-specific rules or human preferences can be more effective and can enable effective maneuvers for fluent human-robot co-navigation.

6 Conclusion

For mobile robots to be safely and efficiently deployed in construction environments, they should be endowed with a certain level of social intelligence to avoid the risk of collisions as well as interruption of construction activities. Our work defined proxemic characteristics relative to work context and incorporated such proxemic requirements into the existing social navigation algorithms so that robots can adapt to the context. We compared the performance of our CASN models against the baseline models. Our preliminary study results showed that the current RL-based socially navigation models can handle low to medium density environments, but they struggled in a high-density environment. This indicates that proxemic-based social navigation algorithms have limitations in the high-density environment, and it is imperative to develop advanced and intricate human-robot interactions to handle such cases. Implicit or explicit human-robot interface can be incorporated to allow robots to more efficiently learn the social norm from human preferences. These findings will help develop and implement mobile robots with social intelligence capable of understanding the workplace context and taking socially and contextually appropriate behaviors accordingly in the construction domain. Such incorporation of the social norm will ensure the psychological and physical safety of workers and support successful integration of mobile robots in the construction sites.

The limitation of our model is that they were performed in a simple simulation environment. It only considers the pedestrian workers, excluding static workers and obstacles. Our model, thereby, may not work well in the real-world environment, which tends to be much complicated. Thus, in future work, we will include both stationary and moving obstacles in addition to different geometric spatial layouts to better represent the characteristics of the real world, and in turn, improve the performance for real-world implementation. Another limitation of our work is that the robot does not infer the social space of the encountering pedestrian workers. Instead, the predetermined values of proxemic requirement were provided to the robot. However, it is unlikely that this value will be fixed or that this knowledge will be provided to robots in real-world scenarios. In our future study, we will make the robot agent capable of inferring the dynamic social space of different workers without explicitly informing them by leveraging vision data collected from robot sensors.

References

[1] X. Shen, E. Marks, N. Pradhananga, T. Cheng, Hazardous Proximity Zone Design for Heavy Construction Excavation Equipment, Journal of Construction Engineering and Management. 142 (2016) 05016001. https://doi.org/10.1061/(ASCE)CO.1943-7862.0001108.

[2] K. Afsari, S. Halder, M. Ensafi, S. DeVito, J. Serdakowski, Fundamentals and Prospects of Four-Legged Robot Application in Construction Progress Monitoring, in: EPiC Series in Built Environment, EasyChair, 2021: pp. 274–283. https://doi.org/10.29007/cdpd.

[3] J.M. Davila Delgado, L. Oyedele, A. Ajayi, L. Akanbi, O. Akinade, M. Bilal, H. Owolabi, Robotics and automated systems in construction: Understanding industry-specific challenges for adoption, Journal of Building Engineering. 26 (2019) 100868. https://doi.org/10.1016/j.jobe.2019.100868.

[4] M.-F.R. Lee, T.-W. Chien, Intelligent Robot for Worker Safety Surveillance: Deep Learning Perception and Visual Navigation, in: 2020 International Conference on Advanced Robotics and Intelligent Systems (ARIS), 2020: pp. 1–6. https://doi.org/10.1109/ARIS50834.2020.9205772.

[5] R. Inoue, T. Arai, Y. Toda, M. Tsujimoto, K. Taniguchi, N. Kubota, Intelligent Control for Illuminance Measurement by an Autonomous Mobile Robot, in: 2019 IEEE International Conference on Advanced Robotics and Its Social Impacts (ARSO), 2019: pp. 270–274. https://doi.org/10.1109/ARSO46408.2019.8948806.

[6] M. Daza, D. Barrios-Aranibar, J. Diaz-Amado, Y. Cardinale, J. Vilasboas, An Approach of Social Navigation Based on Proxemics for Crowded Environments of Humans and Robots, Micromachines. 12 (2021) 193. https://doi.org/10.3390/mi12020193.

[7] T. Kruse, A.K. Pandey, R. Alami, A. Kirsch, Human-aware robot navigation: A survey, Robotics and Autonomous Systems. 61 (2013) 1726–1743. https://doi.org/10.1016/j.robot.2013.05.007.

[8] From Proxemics Theory to Socially-Aware Navigation: A Survey | SpringerLink, (n.d.). https://link.springer.com/article/10.1007/s12369-014-0251-1 (accessed July 22, 2021).

[9] Y. Chen, C. Liu, B.E. Shi, M. Liu, Robot Navigation in Crowds by Graph Convolutional Networks With Attention Learned From Human Gaze, IEEE Robotics and Automation Letters. 5 (2020) 2754–2761. https://doi.org/10.1109/LRA.2020.2972868.

[10] M. Everett, Y.F. Chen, J.P. How, Motion Planning Among Dynamic, Decision-Making Agents with Deep Reinforcement Learning, ArXiv:1805.01956 [Cs]. (2018). http://arxiv.org/abs/1805.01956 (accessed July 28, 2021).

[11] Y.F. Chen, M. Everett, M. Liu, J.P. How, Socially Aware Motion Planning with Deep Reinforcement Learning, ArXiv:1703.08862 [Cs]. (2018). http://arxiv.org/abs/1703.08862 (accessed March 14, 2021).
[12] F. Lindner, C. Eschenbach, Towards a Formalization of Social Spaces for Socially Aware Robots, in: M. Egenhofer, N. Giudice, R. Moratz, M. Worboys (Eds.), Spatial Information Theory, Springer, Berlin, Heidelberg, 2011: pp. 283–303. https://doi.org/10.1007/978-3-642-23196-4_16.
[13] E.T. (Edward T. Hall 1914-2009., The hidden dimension, [1st ed.]., Doubleday, Garden City, N.Y., 1966.
[14] E.T. Hall, R.L. Birdwhistell, B. Bock, P. Bohannan, A.R. Diebold, M. Durbin, M.S. Edmonson, J.L. Fischer, D. Hymes, S.T. Kimball, W. La Barre, S.J. Frank Lynch, J.E. McClellan, D.S. Marshall, G.B. Milner, H.B. Sarles, G.L. Trager, A.P. Vayda, Proxemics [and Comments and Replies], Current Anthropology. 9 (1968) 83–108.
[15] H.I. Moud, I. Flood, C. Capano, Y. Zhang, B. Abbasnejad, Safety of Ground Robot Operations in Construction Job Sites: A Qualitative Approach, in: Construction Research Congress 2020, American Society of Civil Engineers, Tempe, Arizona, 2020: pp. 1327–1335. https://doi.org/10.1061/9780784482865.140.
[16] H. Izadi Moud, I. Flood, X. Zhang, B. Abbasnejad, P. Rahgozar, M. McIntyre, Quantitative Assessment of Proximity Risks Associated with Unmanned Aerial Vehicles in Construction, Journal of Management in Engineering. 37 (2021) 04020095. https://doi.org/10.1061/(ASCE)ME.1943-5479.0000852.
[17]H.I. Moud, I. Flood, A. Shojaei, Y. Zhang, X. Zhang, M. Tadayon, M. Hatami, Qualitative Assessment of Indirect Risks Associated with Unmanned Aerial Vehicle Flights over Construction Job Sites, (2019) 83–89. https://doi.org/10.1061/9780784482445.011.
[18] H. Kretzschmar, M. Spies, C. Sprunk, W. Burgard, Socially compliant mobile robot navigation via inverse reinforcement learning, The International Journal of Robotics Research. 35 (2016) 1289–1307. https://doi.org/10.1177/0278364915619772.
[19] D. Helbing, P. Molnar, Social Force Model for Pedestrian Dynamics, Phys. Rev. E. 51 (1995) 4282–4286. https://doi.org/10.1103/PhysRevE.51.4282.
[20] M. Kollmitz, T. Koller, J. Boedecker, W. Burgard, Learning Human-Aware Robot Navigation from Physical Interaction via Inverse Reinforcement Learning, in: 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2020: pp. 11025–11031. https://doi.org/10.1109/IROS45743.2020.9340865.
[21] M. Luong, C. Pham, Incremental Learning for Autonomous Navigation of Mobile Robots based on Deep Reinforcement Learning, J Intell Robot Syst. 101 (2021) 1. https://doi.org/10.1007/s10846-020-01262-5.
[22] J. Schulman, S. Levine, P. Moritz, M.I. Jordan, P. Abbeel, Trust Region Policy Optimization, ArXiv:1502.05477 [Cs]. (2017). http://arxiv.org/abs/1502.05477 (accessed June 22, 2021).
[23] RVO2 Library - Reciprocal Collision Avoidance for Real-Time Multi-Agent Simulation, RVO2 Library - Reciprocal Collision Avoidance for Real-Time Multi-Agent Simulation. (n.d.). https://gamma.cs.unc.edu/RVO2/ (accessed May 3, 2021).
[24] J. van den Berg, S.J. Guy, M. Lin, D. Manocha, Reciprocal n-Body Collision Avoidance, in: C. Pradalier, R. Siegwart, G. Hirzinger (Eds.), Robotics Research, Springer Berlin Heidelberg, Berlin, Heidelberg, 2011: pp. 3–19. https://doi.org/10.1007/978-3-642-19457-3_1.
[25] K. Li, Y. Xu, J. Wang, M.Q.- Meng, SARL∗: Deep Reinforcement Learning based Human-Aware Navigation for Mobile Robot in Indoor Environments, in: 2019 IEEE International Conference on Robotics and Biomimetics (ROBIO), 2019: pp. 688–694. https://doi.org/10.1109/ROBIO49542.2019.8961764.
[26] C. Chen, Y. Liu, S. Kreiss, A. Alahi, Crowd-Robot Interaction: Crowd-Aware Robot Navigation With Attention-Based Deep Reinforcement Learning, in: 2019 International Conference on Robotics and Automation (ICRA), 2019: pp. 6015–6022. https://doi.org/10.1109/ICRA.2019.8794134.
[27] Robot navigation in dense human crowds: Statistical models and experimental studies of human–robot cooperation - Pete Trautman, Jeremy Ma, Richard M. Murray, Andreas Krause, 2015, (n.d.). https://journals.sagepub.com/doi/abs/10.1177/0278364914557874 (accessed July 22, 2021).