SOCIALGYM 2.0: Simulator for Multi-Robot Learning and Navigation in Shared Human Spaces

Rohan Chandra*, **Zayne Sprague***, **Joydeep Biswas**

The University of Texas at Austin
{rchandra, zaynesprague, joydeepb}@utexas.edu

Abstract

We present SOCIALGYM 2.0, a simulator for multi-agent navigation research. Our simulator enables navigation for multiple autonomous agents, replicating real-world dynamics in complex indoor environments, including doorways, hallways, intersections, and roundabouts. Unlike current simulators that concentrate on single robots in open spaces, SOCIALGYM 2.0 employs multi-agent reinforcement learning (MARL) to develop optimal navigation policies for multiple robots with diverse, dynamic constraints in complex environments. SOCIALGYM 2.0 also departs from the accepted software design standards by employing a configuration-over-convention paradigm providing the capability to benchmark different MARL algorithms, as well as customize observation and reward functions. Users can additionally create their own environments and evaluate various algorithms, based on both deep reinforcement learning as well as classical navigation, using a broad range of social navigation metrics.

Introduction and Related Work

To successfully deploy robots in complex human environments such as airports, grocery stores, hospitals, restaurants, and homes, it is essential to ensure social compliance between robots and humans. But in order to effectively train robots to be socially compliant, we require simulation environments that reflect the complexity of the real world– that is, simulators must provide multi-agent learning support, tightly constrained indoor scenarios, realistic robot and human motion models, and lastly, simulators must be configurable and extensible.

Current simulators only *partially* satisfy the above requirements (Kästner et al. 2022; Tsoi et al. 2022; Chen et al. 2019; Aroor, Esptein, and Korpan 2017; Helbing and Molnar 1995; Biswas et al. 2022; Grzeskowiak et al. 2021; Holtz and Biswas 2021). All of these simulators are currently single-agent navigation in simple open environments. In addition, these simulators model human crowds using reciprocal policies (Van Den Berg et al. 2011) or replay stored trajectories from a dataset (Biswas et al. 2022), or both. Furthermore, while several simulators (Tsoi et al. 2022; Biswas et al. 2022; Grzeskowiak et al. 2021; Holtz and Biswas 2021) model real-world robot dynamics and kinematics realistically, only two simulators (CrowdBOT and our previous work, SocialGym) allow configurability and extensibility to experiment and benchmark different robot kinodynamic configurations. We introduce SOCIALGYM 2.0[1], an open-source simulator for multi-agent social autonomous navigation in challenging environments. SOCIALGYM 2.0 is unique in that it goes–

1. **beyond single agent navigation:** Our multi-agent reinforcement learning paradigm trains multiple autonomous agents each with their own policy. Using the PettingZoo (Terry et al. 2021) and Stable Baselines3 (Raffin et al. 2021) APIs, SOCIALGYM 2.0 supports multi-agent reinforcement learning.
2. **beyond open spaces:** SOCIALGYM 2.0 simulates challenging environments including university campus buildings and geometrically constrained social mini-game scenarios.
3. **beyond convention-over-configuration paradigms:** SOCIALGYM 2.0 uses the *configuration*-over-*convention* paradigm providing users control over every module of the stack, enabling users to conduct research in agent modeling, trajectory planning and collision avoidance, policy learning, and navigation in different social contexts, *simultaneously* keeping the stack simple to use.

The SOCIALGYM 2.0 Design & Architecture

We developed SOCIALGYM 2.0 keeping configurability, extensibility, and modularity in mind, using a configuration-over-convention style. For easy development and research on various aspects of social navigation, we stratified SOCIALGYM 2.0's stack into different layers. At the top of the stack is the PettingZoo (Terry et al. 2021) and Stable Baselines3 (Raffin et al. 2021) interface. This interface uses ROS to send actions from a policy to UTMRS[2], a lightweight simulation engine that acts as an intermediate between the interface and the local navigation and the human crowd simulation modules. The local navigation planner is responsible for

* denotes equal contribution

[1]The full paper, code, and documentation can be found at https://amrl.cs.utexas.edu/SocialGym2/index.html.

[2]University of Texas Multi-Robot Simulator

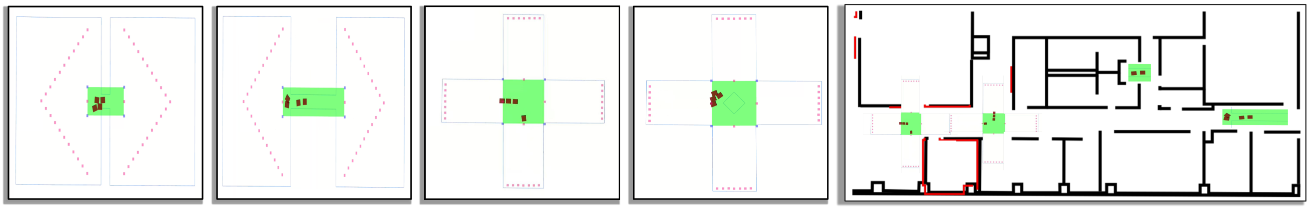

(a) Doorway (b) Hallway (c) Intersection (d) Roundabout (e) Multiple scenarios

Figure 1: SOCIALGYM 2.0 includes five types of social environments–Doorway, Hallway, Intersection, Roundabout, and Multi-Scenario consisting of multiple intersection, doorway, and hallway scenarios.

converting high-level actions from the PettingZoo interface into continuous motion commands that satisfy the underlying robot dynamics and sends back the next state to the simulation engine. Each layer of the stack has a modular API that allows researchers and developers to focus on a single part of the stack at a time without having to refactor or access other parts of the stack.

Training and Evaluating a Navigation Policy

The life-cycle of training a multi-agent policy in the interface mimics that of OpenAI Gym or Stable-Baselines. It begins with selecting 2D maps and scenarios. Then, two class definitions, the Observer and Rewarder, are used for tracking the observations and rewards during each step. The environment is then defined, which instantiates the Gym Environment and initializes the ROS submodules with all the information needed to load the first 2D map and scenario. Next, a Stable Baselines-v3 Policy is chosen and the learning method can be called to train the policy. Finally, the policy is evaluated on metrics for socially compliant navigation.

2D Maps, Navigation Graphs, and Scenarios

We developed a Docker-based program for creating 2D maps. These maps consist of vectors representing obstacles and a navigation graph that defines possible paths. The maps enable the creation of "social mini-games" for training and evaluating agents in challenging situations. Additionally, we can use 2D floor plans for larger, more realistic scenarios.

Our simulator, SOCIALGYM 2.0, provides pre-defined indoor scenarios and allows users to easily extend them by adding obstacles and walls. It supports multiple social mini-games, and parameters like door/hallway width and conflict zones can be adjusted. Agent density and kinodynamic constraints can also be modified, resulting in heterogeneous behavior. Simulated pedestrians can be customized to add further complexity.

Observations and Rewards

The complexity of social navigation lies in the various definitions and representations used, with no universal metric available (Mavrogiannis et al. 2021; Mirsky et al. 2021). To address this, we offer full customization of the state space and reward functions in our open-source simulator. Researchers can define their own notions of social navigation by making minimal changes to the code. We provide a flexible class structure that allows for the creation of custom observations and reward functions, expanding the capabilities of SOCIALGYM 2.0.

Evaluation Metrics

In addition to the standard evaluation metrics available in Stable Baselines3 and PettingZoo, we extend these functions and implement custom evaluation metrics in the same style of SocNavBench and SEAN2.0. This stems from the noted ambigious definition of social navigation. Often, in lieu of a single metric that best defines Social Compliance, multiple metrics are used as a proxy. We implemented the most common metrics used to measure social compliance, including partial and full success rates, velocity changes, average stopping time, collision rates, and more.

Social Mini-game Benchmarks

In this work, we include five mini-game scenarios to benchmark social navigation. These are ODoorway, Hallway, Intersection, Roundabout, and Multi-Scenario, depicted in Figure 1. Social mini-games may be described as a scenario with multiple agents accomplishing a shared goal in spatially constrained environments. Such scenarios frequently arise indoors at schools, hospitals, airports, etc. as well as outdoors on sidewalks and traffic intersections. We provide a point-click/drag interface that allows for the easy construction of both vector maps and navigation graphs.

Conclusion, Limitations, and Future Work

In conclusion, this paper presents SOCIALGYM 2.0, a multi-agent navigation simulator for social navigation research. SOCIALGYM 2.0 has certain limitations that are currently being addressed, including constraints on CPU resources and lack of optimization for multi-threading and parallel processing. We also plan to enhance practicality by introducing continuous actions and state-of-the-art MARL algorithms such as QMIX, MAPPO, and MADDPG.

Acknowledgments

This work is partially supported by the National Science Foundation (CAREER-2046955 and CCF-2006404). Any opinions, findings, and conclusions expressed in this material are those of the authors and do not necessarily reflect the views of the sponsors.

References

Aroor, A.; Esptein, S. L.; and Korpan, R. 2017. Mengeros: A crowd simulation tool for autonomous robot navigation. In *2017 AAAI Fall Symposium Series*.

Biswas, A.; Wang, A.; Silvera, G.; Steinfeld, A.; and Admoni, H. 2022. Socnavbench: A grounded simulation testing framework for evaluating social navigation. *ACM Transactions on Human-Robot Interaction (THRI)*, 11(3): 1–24.

Chen, C.; Liu, Y.; Kreiss, S.; and Alahi, A. 2019. Crowd-robot interaction: Crowd-aware robot navigation with attention-based deep reinforcement learning. In *2019 International Conference on Robotics and Automation (ICRA)*, 6015–6022. IEEE.

Grzeskowiak, F.; Gonon, D.; Dugas, D.; Paez-Granados, D.; Chung, J. J.; Nieto, J.; Siegwart, R.; Billard, A.; Babel, M.; and Pettré, J. 2021. Crowd against the machine: A simulation-based benchmark tool to evaluate and compare robot capabilities to navigate a human crowd. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 3879–3885. IEEE.

Helbing, D.; and Molnar, P. 1995. Social force model for pedestrian dynamics. *Physical review E*, 51(5): 4282.

Holtz, J.; and Biswas, J. 2021. SOCIALGYM: A Framework for Benchmarking Social Robot Navigation. *arXiv preprint arXiv:2109.11011*.

Kästner, L.; Bhuiyan, T.; Le, T. A.; Treis, E.; Cox, J.; Meinardus, B.; Kmiecik, J.; Carstens, R.; Pichel, D.; Fatloun, B.; et al. 2022. Arena-bench: A benchmarking suite for obstacle avoidance approaches in highly dynamic environments. *IEEE Robotics and Automation Letters*, 7(4): 9477–9484.

Mavrogiannis, C.; Baldini, F.; Wang, A.; Zhao, D.; Trautman, P.; Steinfeld, A.; and Oh, J. 2021. Core challenges of social robot navigation: A survey. *arXiv preprint arXiv:2103.05668*.

Mirsky, R.; Xiao, X.; Hart, J.; and Stone, P. 2021. Prevention and resolution of conflicts in social navigation–a survey. *arXiv preprint arXiv:2106.12113*.

Raffin, A.; Hill, A.; Gleave, A.; Kanervisto, A.; Ernestus, M.; and Dormann, N. 2021. Stable-Baselines3: Reliable Reinforcement Learning Implementations. *Journal of Machine Learning Research*, 22(268): 1–8.

Terry, J.; Black, B.; Grammel, N.; Jayakumar, M.; Hari, A.; Sullivan, R.; Santos, L. S.; Dieffendahl, C.; Horsch, C.; Perez-Vicente, R.; et al. 2021. Pettingzoo: Gym for multi-agent reinforcement learning. *Advances in Neural Information Processing Systems*, 34: 15032–15043.

Tsoi, N.; Xiang, A.; Yu, P.; Sohn, S. S.; Schwartz, G.; Ramesh, S.; Hussein, M.; Gupta, A. W.; Kapadia, M.; and Vázquez, M. 2022. Sean 2.0: Formalizing and generating social situations for robot navigation. *IEEE Robotics and Automation Letters*, 7(4): 11047–11054.

Van Den Berg, J.; Guy, S. J.; Lin, M.; and Manocha, D. 2011. Reciprocal n-body collision avoidance. In *Robotics Research: The 14th International Symposium ISRR*, 3–19. Springer.