RL-Godot Pedestrian Simulation: Curriculum–Based Reinforcement Learning for Pedestrian Simulation

Giuseppe Vizzari*[1], Andrea Falbo[1], Ruben Tenderini[1], and Daniela Briola[1]

[1]University of Milano-Bicocca, Italy

Abstract The abstract describes a research effort aimed at developing a Reinforcement Learning based open source software system supporting the study of pedestrian and crowd dynamics. We briefly discuss the system and a first experimental application aimed at evaluating its adequacy to simulate medium to high local density situations.

Keywords Reinforcement Learning, Curriculum Learning, Pedestrian simulation

Extended abstract

Reinforcement Learning (RL) [1] is a machine learning approach in which an agent learns how to behave in a given environment in which it is situated, thanks to a feedback signal provided by the environment itself as a consequence of an action carried out by the agent. This feedback, a scalar reward signal, is much less informative than the loss function value typical of neural network training. Nonetheless, it is applicable in cases where there are difficulties in achieving sufficiently large and varied data sets for supervised training, but an evaluation of the quality of an action can be computed.

Pedestrian simulation, despite being a very active field of research, has provided empirical evidence allowing for such an evaluation. Recent works have shown that Deep Reinforcement Learning (DRL) (i.e. RL techniques employing deep neural networks in their functioning) can produce policies producing plausible pedestrian behaviors. In particular, [2] employed a system based on Unity and the associated ML-Agents toolkit[1], as technical enabling technologies, coupled with a curriculum learning approach, to achieve a general model that produces plausible pedestrian behaviors in environments not approached during training. Employing a game engine allows both performing simulations and exporting data for external analyses, but it also allows employing the achieved models within games, interactive applications, even in Vitual Reality experiments for acquiring empirical data about pedestrian behavior.

Unity, however, is not an open source platform, and ML-Agents does not represent a generalized library providing a wide range of DRL algorithms, but rather an ancillary complement to Unity. We decided therefore to develop an analogous platform only employing open source technologies and in particular: (i) the **Godot** [2] game engine, (ii) the **Godot RL Agents** library [3] that grants the possibility to connect Godot to external RL and DRL libraries such as **Stable Baselines** [4], an open source library providing a set of reliable implementations of a large number of DRL in PyTorch; data about the simulated scenarios can then be plotted employing **PedPy**[5], an open source e library for pedestrian movement analysis. The result is RL-Godot Pedestrian Simulation [6], a system granting the possibility to perform simulations in an environment created respecting specific requirements employing some already available agent behavioral models, or even customize the training process to generate new behavioral models.

First of all, we reproduced the results of the above-mentioned work [2], then we considered situations characterized by a higher level of pedestrian density, such as the scenarios described in [3] and [4]: in this kind of situation, the model achieved through the previous work, and the enclosed training process, did not lead to plausible results. Our conjecture was that the reward function was excessively penalizing agents for getting close to other ones. We performed a new training process, reducing the penalization for the invasion of proxemic space, as well as including in the training curriculum additional scenarios in which learning agents encountered higher density situations.

*Email of the corresponding author: `giuseppe.vizzari@unimib.it`
[1]`https://github.com/Unity-Technologies/ml-agents`
[2]`https://godotengine.org/`
[3]`https://github.com/edbeeching/godot_rl_agents`
[4]`https://github.com/DLR-RM/stable-baselines3`
[5]`https://pedpy.readthedocs.io/`
[6]`https://github.com/Ruben-2828/RL-Godot-Pedestrian-Simulation`

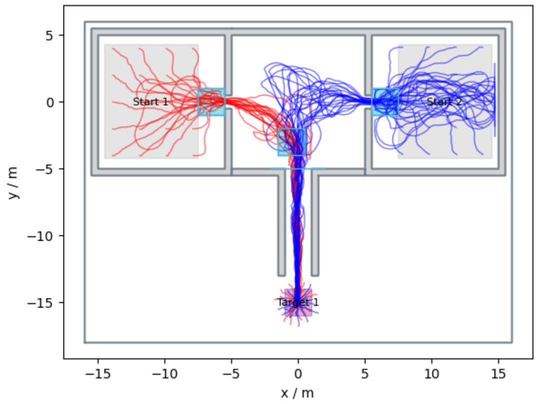

(a) Trajectories generated by agents in the center door connection

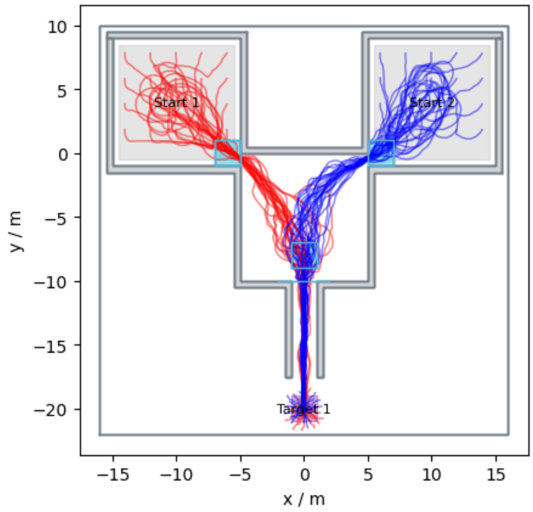

(b) Trajectories generated by agents in the corner door connection

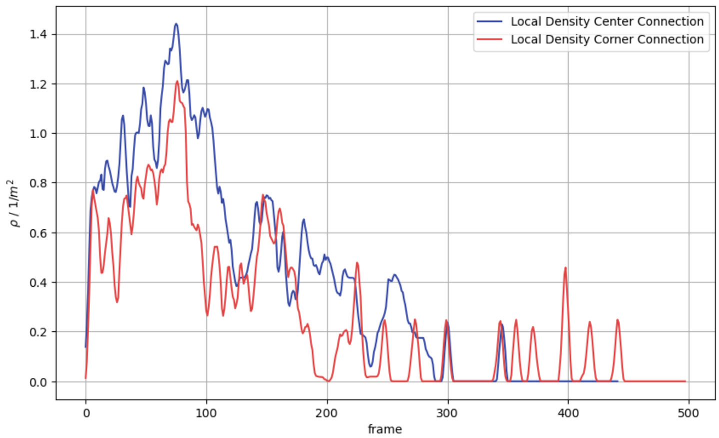

(c) Overall measured density in the central room

Figure 1: Results achieved in the scenario discussed in [3].

The achieved pedestrian behavioral model performed significantly better than the previous one: Fig. 1(a) and Fig. 1(b) respectively show trajectories in the two considered spatial configurations of the environment, and Fig. 1(c) show the level of density in the monitored area (highlighted as a cyan square in the previous figures) throughout the simulation. Results are encouraging and comparable to those reported in the cited papers.

Bibliography

[1] Sutton, R.S. and Barto, A.G., 1998. Reinforcement learning: An introduction (Vol. 1, No. 1, pp. 9-11). Cambridge: MIT press.

[2] Vizzari, G. and Cecconello, T., 2022. Pedestrian simulation with reinforcement learning: A curriculum-based approach. Future Internet, 15(1), p.12.

[3] Ezaki, T., Yanagisawa, D. and Nishinari, K., 2012. Pedestrian flow through multiple bottlenecks. Physical Review E—Statistical, Nonlinear, and Soft Matter Physics, 86(2), p.026118.

[4] Adrian, J., Seyfried, A. and Sieben, A., 2020. Crowds in front of bottlenecks at entrances from the perspective of physics and social psychology. Journal of the Royal Society Interface, 17(165), p.20190871.