ROB 538 Project: Dynamic Robot Barriers

ALEXANDRA BACULA, RHIAN PRESTON, and COLIN SHEA BLYMYER

In places with a lot of people in one space, crowd control is often an issue that is taken care of with humans controlling traffic flow or with static barriers. Human traffic controllers can adapt to changing situations, such as an evacuation or sudden influx of people during rush hour, but it poses a safety hazard for them. Multi-agent robot systems provide a novel approach to flexible human traffic flow control for both day-to-day traffic and emergency scenarios. Modern human traffic simulation approaches and reinforcement learning are used as a test environment for determining robot behaviors and roles for changing human traffic flow and a dynamic way. Results show that such a complex environment poses challenges for classical reinforcement learning approaches, but such challenges may be overcome with more advanced techniques.

1 INTRODUCTION

In situations where crowds gather, crowd control is necessary for mitigating regular variances in congestion and even more critical in effective evacuation during emergency scenarios. Maintaining the flow of pedestrian traffic in emergency scenarios is vital to public safety. However, the static traffic control measures common in places like subway platforms are not well equipped for adapting to drastic changes in traffic flow, and employing humans to control traffic exposes them to the very dangers the traffic is fleeing.

Crowd dynamics are complex, exhibiting a variety of behavior patterns as a homogeneous system, as well as being composed of discrete actors with independent thoughts and actions. These dynamics are a research field in and of themselves, and any crowd control system must make choices and take actions while understanding the effect those actions will have on the crowd dynamics. A poorly designed traffic control system that does not account for these behaviors risks causing further traffic problems, rather than improving traffic flow.

A multi-robot solution could provide dynamic reconfiguration of the traffic space depending on the situation at hand as compared to static barriers. Such a system could improve response time in an emergency situation, as compared to the time required to deploy human responders. However, this approach is not without challenges. In addition to the complexity of a crowd environment, multi-agent systems introduce complexities regarding inter-agent communication and learning.

In this work we introduce two transient crowd flow scenarios and compare the effectiveness of dynamic and static control measures in these scenarios. We implement these scenarios as simulations to show the effectiveness of dynamic crowd flow control in changing traffic patterns. We first review related work, then go into details of our solution and implementation, and finally our results and conclusions.

2 RELATED WORK

In order to fully understand the domain we are working in, we look at four areas of related work. We start with multi-robot group formation and behavior, as we plan to expand on these methods to create our own multi-robot algorithms. Next we look at social navigation, to see how robots currently navigate in human spaces. Then we examine human crowd dynamics, which are necessary to understand so that we can use robots to improve traffic flow. Finally we look at simulations of human crowds as we plan to modify an existing simulation to test our work.

Authors' address: Alexandra Bacula, baculaa@oregonstate.edu; Rhian Preston, prestonr@oregonstate.edu; Colin Shea Blymyer, sheablyc@oregonstate.edu.

2.1 Multi-robot group formation and behavior

The first area of interest is multi-robot group behavior. For this project, the focus will be on distributed group behavior methods, as in the real world, it is unlikely that the robots will have a centralized system. One type of distributed control uses behavioral models, such as flocking. Flocking involves all robots following a defined set of individual actions that lead to an overall group behavior [2]. Early work in this area started in animation. In work by Reynolds, a flock of birds is animated by modelling the birds as particles who react in certain ways to their local environment, instead of creating a pre-planned trajectory for each bird [20]. This behavior based approach was extended to robotics, with one application being unmanned ground vehicles working as a team to reach goals and perform tasks [1]. One example of behavior rules bases the individual behavior on the heading of the robot and the headings of the robots adjacent to it [9].

2.2 Social Navigation

Next, we examine social navigation to see how previous work has tackled the issue of robot navigation in a human friendly way. One method for socially cognizant navigation is potential field based planning [3, 11]. This approach involves creating different social “forces" that the robot is influenced by to drive how the robot gets from one place to another. This may be used as a global path planner, or as a local planner in combination with another planning method for the global planning, such as A*. Work by Kruse et al. [11] takes into account that humans have agency and are not just obstacles, and that the quality of the plan depends on the human acceptance more than its effectiveness. Ferrer took inspiration from Helbing’s social force model of humans [8] and applied it to the context of a companion robot that walks alongside a person [3, 4]. Converting human social rules into cost functions and maps for planners is another popular approach for integrating human social proxemics into motion planning. Kirby turned human navigation rules into cost functions for a heuristic planner for a single robot that resulted in person-acceptable navigation [10]. Simulations were run using the planner which validated that socially acceptable navigation choices were made by the robot in different contexts.

Another method for social navigation is learning. Vasquez et al. [22] used inverse reinforcement learning to model action motivation and context into learning. They used the ROS pedestrian simulator as a testbed and evaluation metrics that were objective and subjective, such as human proxemics and an influence model based on work by [7]. Okal and Arras [17] also used inverse reinforcement learning to learn normative social structures for humans and create trajectories to accommodate the humans. Results showed the robot was able to learn acceptable trajectories in both large scale pedestrian simulation and a small in-person study [17]. Luber et al. [13] took a different approach and used unsupervised learning on annotated surveillance data to generate relative motion paths that well replicated human relative motion.

2.3 Human Crowd Dynamics

Another area of interest is the dynamics of crowds and how they are currently influenced. Work by Helbing discusses the concept of social force, defined by an individual’s motivation, and how it can be used in creating a model for pedestrian behavior [8]. The social force model assumes that a person will go the most direct route to their destination, that each individual pedestrian will want to keep some distance between themselves and others and that this distance is a function of speed. The social force equation then consists of an attractive force towards the goal and a repulsive force towards others. Simulations done showed that the social force model was able to decently reproduce pedestrian movement. Additional work by Helbing studied human traffic flow in different environments including bottlenecks, intersections, and corridors, with specific focus on how these features can effect evacuation of a space [6]. These simulations showed that the geometry of the space effects how people exit the space, and that additions such as a zig-zag path or columns can help reduce panic and congestion in crowds. Previous simulations and real life experiments have shown that changing the environment in specific ways, such as adding pillars or obstacles to a space, can aid in making evacuation more efficient, as seen in Figures 1 and 2 [6, 14].

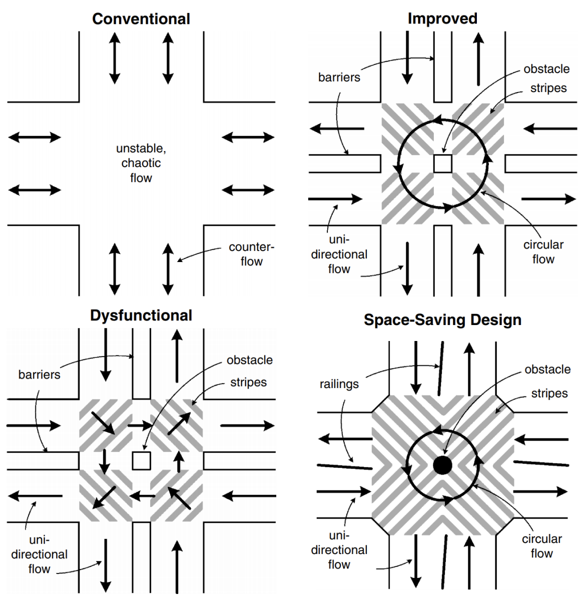

Fig. 1. Intersections modified with additional barriers, railings, obstacles, and visual signals to aid in optimizing human traffic flow [6].

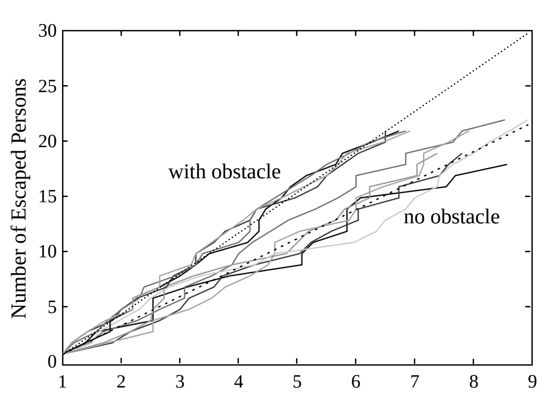

Fig. 2. A real-life experiment involving a panicked evacuation scenario showed that particular placement of an obstacle can aid in the efficiency of evacuation [6].

2.4 Simulations of Human Crowds

In order to properly analyze any multi-robot solution, it must be tested within appropriate crowd scenarios. Simulating these crowds and scenarios with dynamic environment changes is a multi-agent system itself with its own body of research work and can be very computationally intensive. Prior work in crowd simulation, both for day-to-day traffic scenarios and more urgent evacuation scenarios includes a variety of methods for creating and bench-marking these simulations. In order to mimic both crowd dynamics, as well as the more localized independent agent behaviors, a standardized approach is to utilize a large, particle-style controller for the crowd, and localized path planning to handle collision avoidance between agents and between agents and obstacles [18, 19, 23]. Work by [18] and [21] focus on creating efficient collision avoidance and reasonable behavior, but both lean on a statically defined environment to determine behaviors.

Work by [23] uses roles to help drive behaviors from a cognitive framework. Leader and follower are the two roles, allowing flexible environments by focusing path planning and collision handling more heavily on the leader agents. In contrast, [16] relies on very localized path planning and collision avoidance, paired with the concept of incompressibility of agents to create behaviors for very dense crowds. This efficiently handles collision avoidance while limiting the needs of computing power. This approach, however, does not account for distant agents for planning, which leads to some non-realistic motion and agent spreading. Further exploration of a cognitive model for high density crowds was performed by [19]. This work provides agents with individual personalities and rule sets, leading to unique behaviors based on agent types. This work functions in highly dynamic environments.

3 METHODS

In order to influence human crowd behavior effectively, the robots create dynamic barriers to allow them to reach their desired configuration and guide people on what to do. In a panic situation, movement is more useful than lights or sound, as it would be likely that sound can be drowned out by alarms and lights can add to the panic. The robots modulate their formations and movements based on the crowd density and other factors (such as rush hour) to create optimal spaces for pedestrian traffic.

For the purposes of this project, we focus only on planar, 2D movement for both the people and robots. The robots will be simulated as differential drive mobile robots with degrees of freedom in x, y, and θ. Our plan is to build on previously created crowd simulation models in Python, namely PySocialForce[5], to create different human traffic simulation environments. We will build different robot barriers and behaviors, as well as look at localized modes and sensing and communication for adaptable positioning. The crowd response and response times to different barriers, behaviors, and changing modes will be examined. Additionally, we will examine the flexibility and speed of formation for the different behaviors, and use the data from the crowd and robots to determine the most promising of the barrier-behavior-communication combinations.

3.1 Crowd Behavior

The crowd behavior used for simulation is modeled using the PySocialForce implementation of the Extended Social Force Model[5, 15]. This implementation extends the base Social Force Model [8] to simulate pedestrian social group walking behaviors. The base implementation uses only static obstacles, and so the package is modified to allow for dynamic obstacles that utilize their own learning method.

3.2 Learning

The dynamic barrier system will include four agents, where each agent is one meter long. Two agents are rotated 45° clockwise, and the other two are rotated 45° counter-clockwise. Each agent has knows where all the pedestrians are, where the other robots are, and the general destinations of each of the pedestrians. The agents do not know anything about the angle of rotation of the other agents or the size of the agents. Each agent can move one meter backwards, forwards or not move in both the x and y planes, leading to nine possible actions.

The agents will employ Q-Learning to minimize the time pedestrians spend traversing the interchange in each scenario. To facilitate learning, the ten meter square interchange is sub-sectioned into one meter bins. These location bins are then further binned based on the number of pedestrians present, as well as if there is already an agent in that particular bin. The population bins are broken into five possible densities, and each density state has an equivalent state that includes the presence of an agent. This breakdown of the space leads to roughly 11^{100} possible states. This estimate does not account for the fact that only four of the 25 states can have agents at any given time, and it is quite possible a significant number of these states will never be observed.

The crowd simulation model does not allow for dynamic generation of pedestrians or continuous flow, so training was done by generating 50 second "snapshot" simulations of the pedestrian flow. To accelerate the learning process with these snapshots, we run six snapshots in parallel at a time, and then average the Q-values across the snapshots for each of the four agents. These averaged Q-values are then used as starting Q-values for the next learning cycle.

In each step, all agents are rewarded with the number of people who have reached the exit at the end of that step.

3.3 Scenarios

The conventional interchange shown in Figure 1 is used as the simulation layout for ease of comparison to previous work. The space is modeled after a pedestrian interchange in a subway system and has four points of ingress and four points of egress.

We investigate two scenarios in which we compare the performance of our dynamic barriers to the performance of static barriers and that of no barriers. Performance is measured by the average flow rate of the environment.

3.3.1 Rush Hour. In the rush hour scenario, barriers must reduce congestion at the intersection of traffic flows from four pairs of ingress and egress points. The layout of this scenario can be seen in Figure 3. Foot traffic will begin slowly and the influx of human agents will increase sharply over time. Human agents will move from an assigned ingress point to an assigned egress point, each starting with a velocity sampled from a normal distribution centered on 1.4 m/s (a standard walking speed [12]). Barriers in this scenario must efficiently direct low and high volumes of traffic to optimize traffic flow. For our rush hour intersection, our static barriers will be a center line barrier in each aisle with a column in the center of the intersection, as seen in the top right of Figure 4 based on the results of Helbing et al. [6]. This configuration is efficient for high densities, which is appropriate for a rush hour situation [6].

3.3.2 Evacuation. In the evacuation scenario, many human agents enter the environment through one of three ingress points with a very high velocity (around 2.8 m/s) with one point of egress. Our base environment can be seen in Figure 5 This will likely form congestion at the egress point, which the barrier system must try to disrupt. As the environment is also an intersection, we chose to have the same barriers as the rush hour scenario based on prior work by Helbing et al. [7]. Our static barrier configuration can be seen in Figure 6.

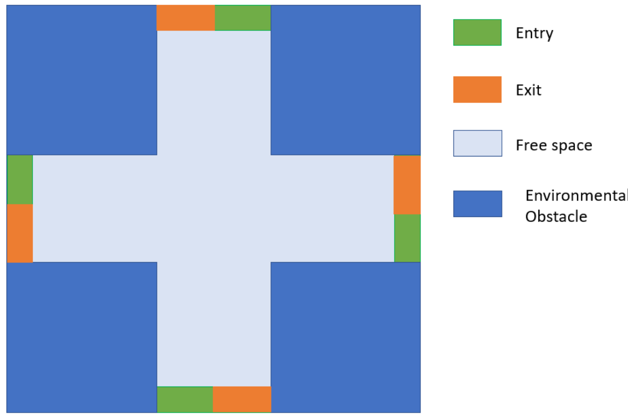

Fig. 3. Environment for a rush hour situation with four points of ingress and egress.

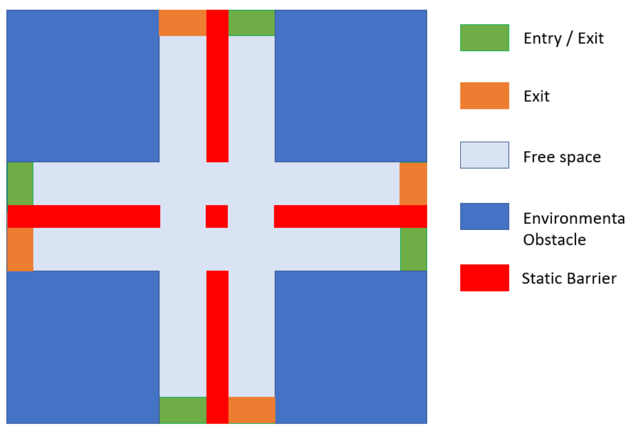

Fig. 4. Environment with static barriers for a rush hour situation with four points of ingress and egress.

3.4 Evaluations

Once learning is complete, we will run two cases for each scenario. The cases for each scenario will use the same generated number of pedestrians and time steps. The first case uses static barriers defined based on previous literature on pedestrian flow. The second case uses four dynamic barriers, where two are cantilevered right and two are cantilevered to the left. Additionally, the q-tables developed during the learning phase are used for these evaluation runs.

4 RESULTS

We will evaluate our system performance with the time to exit of 95% of people in the environment. We are attempting to minimize the time to exit, and we aim for the time metric to be better than the control experiment. In both scenarios, we expect to see that there will be an improved performance using dynamic barriers over static ones in most cases.

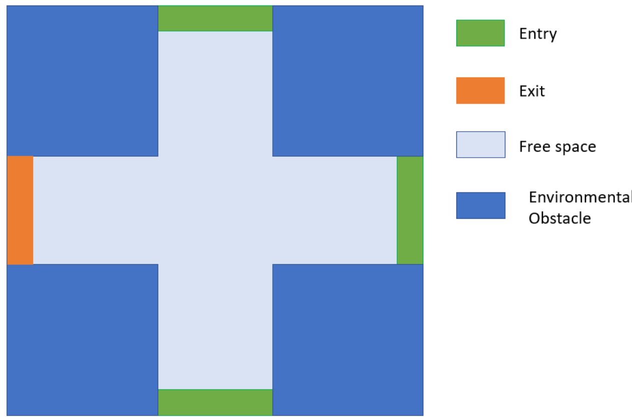

Fig. 5. Environment for an evacuation situation with one point of egress.

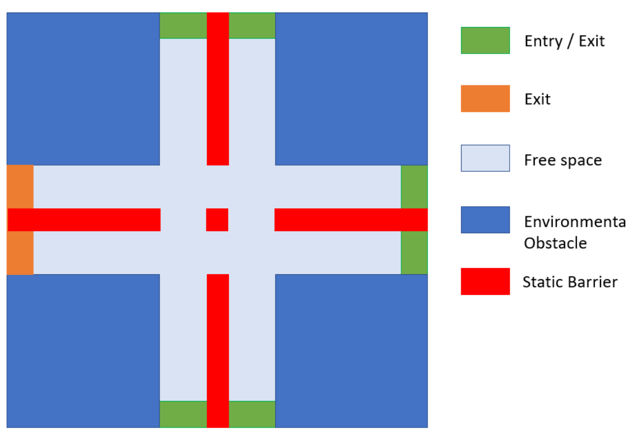

Fig. 6. Environment with static barrier for an evacuation situation with one point of egress.

4.1 Rush Hour

We hypothesize that static barriers will improve the time to exit, as seen in prior work outline in previous sections. We hypothesize that dynamic barriers will help improve time to exit in a rush hour scenario, but not as significantly as in an evacuation scenario, as new people are entering the space more frequently, keeping the density of people more constant than the evacuation scenario. The less constrained setup of the rush hour scenario may make learning harder than the evacuation scenario. The barriers are also expected to be less dynamic than the evacuation scenario as the rush hour scenario changes at a slower pace.

4.2 Evacuation

We hypothesize that static barriers will improve the time to exit, as seen in prior work outlined in previous sections. We hypothesize that dynamic barriers will help more in an evacuation scenario than a rush hour scenario, as the number of people is constantly decreasing and no new people are entering the space. The more constrained setup of this scenario may make learning easier, yielding

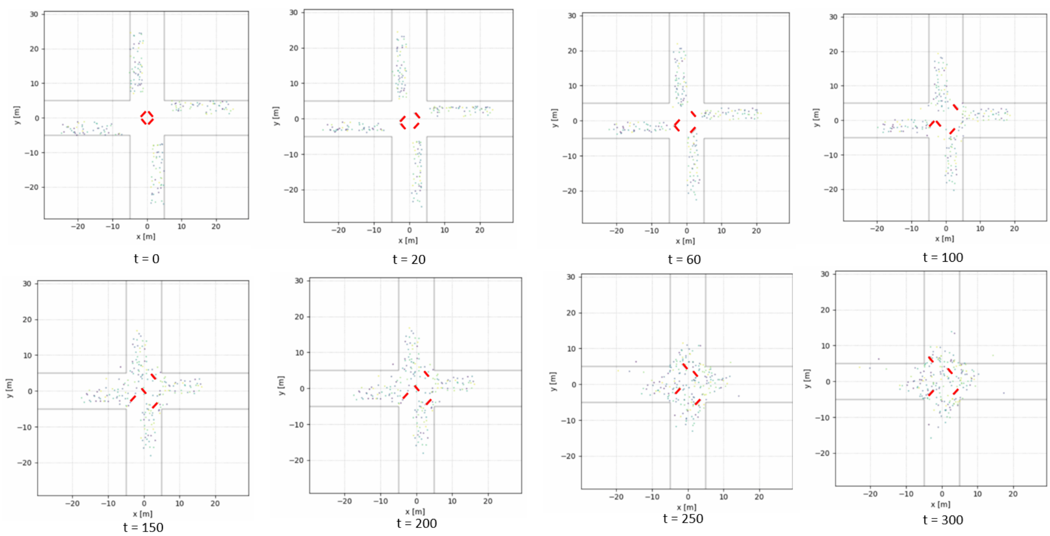

Fig. 7. Dynamic barriers in rush hour scenario. Red lines are dynamic barriers

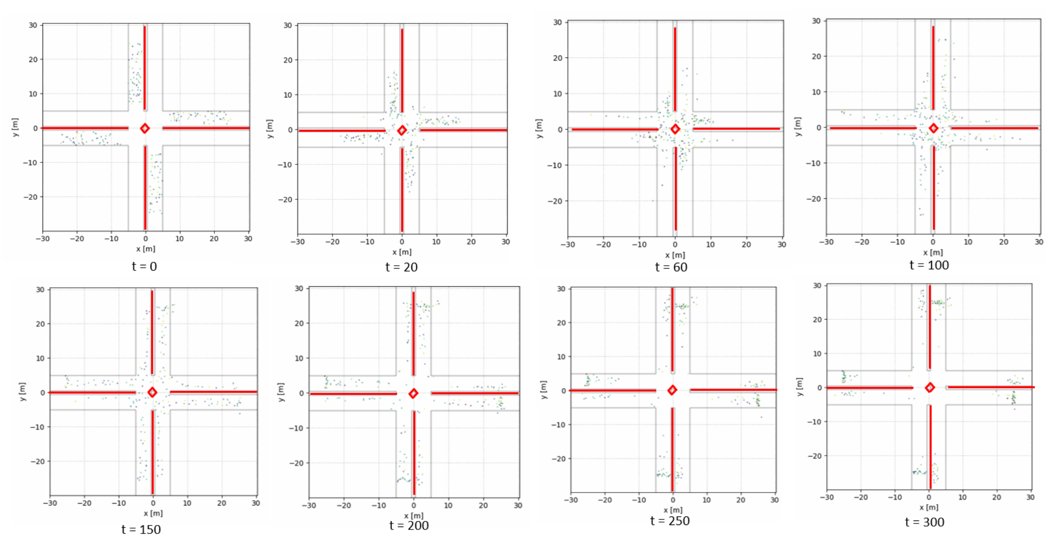

Fig. 8. Static barriers in rush hour scenario. Red lines are dynamic barriers

better results. The barriers are also expected to be more dynamic in the evacuation scenario as it is a faster paced scenario with the number of people changing quicker.

5 CONCLUSION

In this work we introduced the concept of dynamic robot barriers for increasing safety and efficacy in human traffic control. Using a human crowd simulator, we used reinforcement learning to train dynamic line barriers to increase traffic flow in an rush hour and evacuation scenario. Results showed that, even for experiments with run-times longer than twelve hours, agents failed to

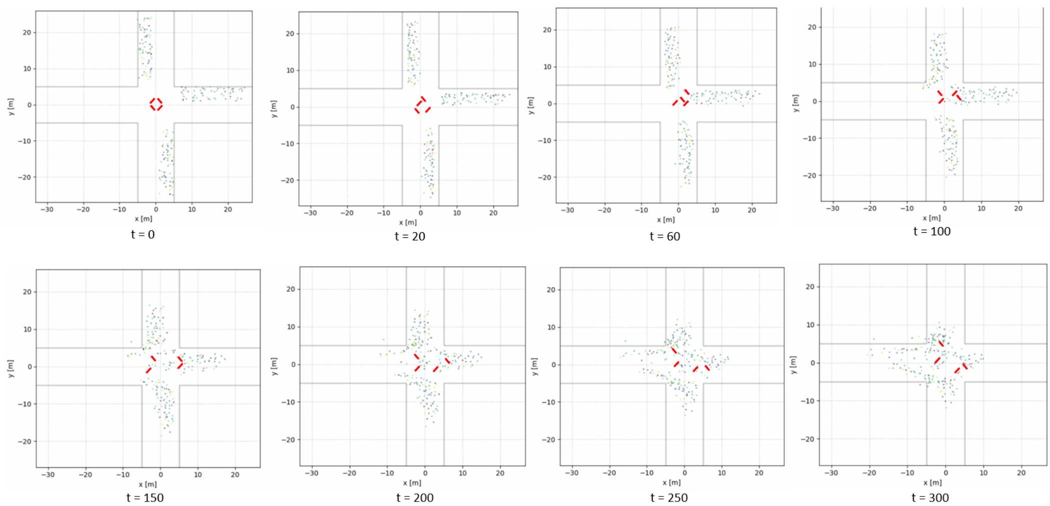

Fig. 9. Dynamic barriers in evacuation scenario. Red lines are dynamic barriers

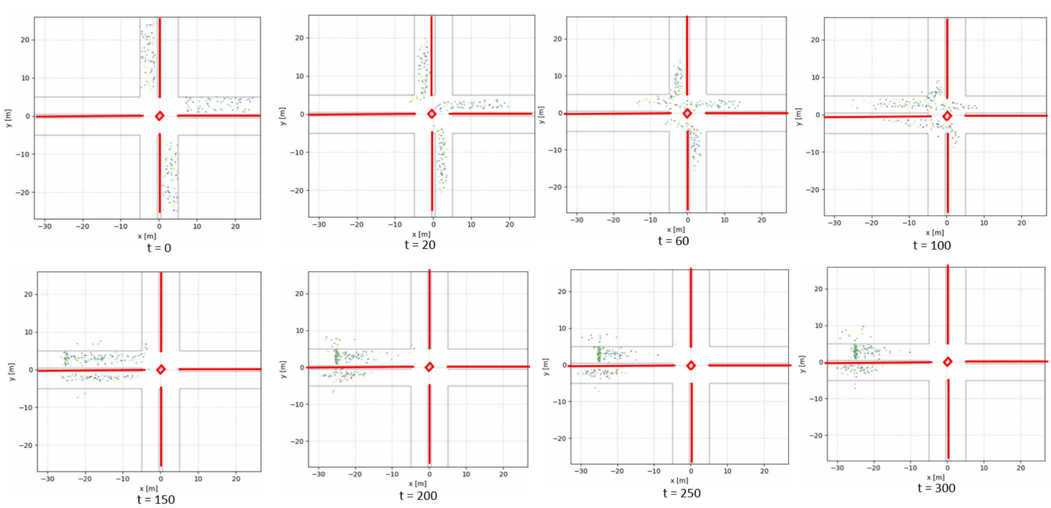

Fig. 10. Static barriers in evacuation scenario. Red lines are dynamic barriers

converge on a policy. We suspect this is a consequence of the large state-action space exposed to the agents, even after quantization and tiling.

Despite these setbacks, our dynamic barriers seem to perform on-par with the static barriers — taking only a few more seconds on average to clear the intersection during a given scenario. This may imply that the environment itself is insensitive to interventions, which would make it difficult to engineer a reward for the agents that improves the throughput of the intersection. Further, the limitations we imposed on the movement of our agents reduces the amount of influence they can impose on the environment.

One goal of future work is to examine use of a neural network for learning, as opposed to classical Q-learning. This would allow for the more complete state space definition, as well as allowing for

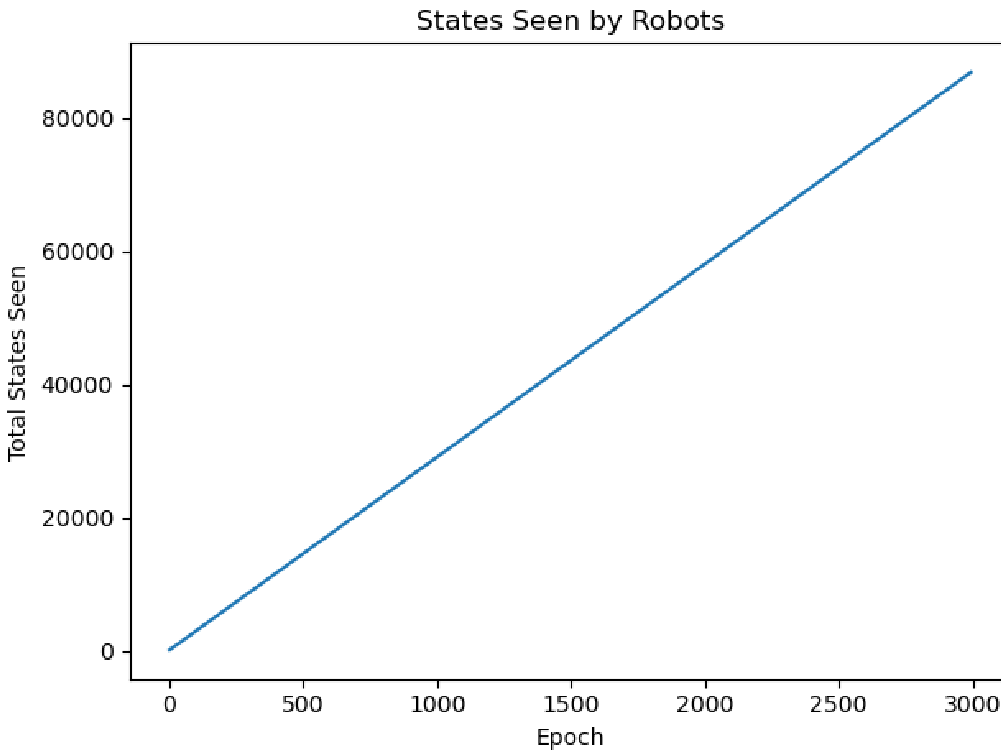

Fig. 11. In the evacuation scenario, each epoch exposed many new states to the agents, even after 3000 epochs. Effective learning only begins happening once the agents have seen most of the states.

more control over the barriers so they can rotate as well as translate. In conjunction with this, it would be interesting try different barrier shapes, for example, changing barrier lengths or including circular barriers in addition to line barriers. Another avenue of future work is changing the rate of people entering the simulation so we can more accurately replicate a rush hour situation to see how the barriers change from light traffic to very heavy traffic. In addition to changing the people flow, another avenue of future work is testing in different environments, such as a single hallway, a room with only one exit, or a more complex intersection.

ACKNOWLEDGMENTS

The contributions of the team members can be seen in Table 1.

Category	Alex	Rhian	Colin
Writing	60%	20%	20%
Coding	0%	60%	40%
Technical	33%	33%	33%
Organization	33%	33%	33%

Table 1. Contribution of Group Members

REFERENCES

[1] Tucker Balch and Ronald C Arkin. 1998. Behavior-based formation control for multirobot teams. *IEEE transactions on robotics and automation* 14, 6 (1998), 926–939.

[2] Y. Cao, W. Yu, W. Ren, and G. Chen. 2013. An Overview of Recent Progress in the Study of Distributed Multi-Agent Coordination. *IEEE Transactions on Industrial Informatics* 9, 1 (Feb 2013), 427–438.

[3] Gonzalo Ferrer, Anais Garrell, and Alberto Sanfeliu. 2013. Robot companion: A social-force based approach with human awareness-navigation in crowded environments. In *Intelligent robots and systems (IROS), 2013 IEEE/RSJ international conference on*. IEEE, 1688–1694.

[4] Gonzalo Ferrer, Anaís Garrell Zulueta, Fernando Herrero Cotarelo, and Alberto Sanfeliu. 2017. Robot social-aware navigation framework to accompany people walking side-by-side. *Autonomous robots* 41, 4 (2017), 775–793.

[5] Yuxiang Gao. [n.d.]. yuxiang-gao/PySocialForce. https://github.com/yuxiang-gao/PySocialForce

[6] Dirk Helbing, Lubos Buzna, Anders Johansson, and Torsten Werner. 2005. Self-organized pedestrian crowd dynamics: Experiments, simulations, and design solutions. *Transportation science* 39, 1 (2005), 1–24.

[7] Dirk Helbing, Illés J Farkas, and Tamás Vicsek. 2002. Crowd disasters and simulation of panic situations. In *The Science of Disasters*. Springer, 330–350.

[8] Dirk Helbing and Peter Molnar. 1995. Social force model for pedestrian dynamics. *Physical review E* 51, 5 (1995), 4282.

[9] Ali Jadbabaie, Jie Lin, and A Stephen Morse. 2003. Coordination of groups of mobile autonomous agents using nearest neighbor rules. *Departmental Papers (ESE)* (2003), 29.

[10] Rachel Kirby, Reid Simmons, and Jodi Forlizzi. 2009. Companion: A constraint-optimizing method for person-acceptable navigation. In *Robot and Human Interactive Communication, 2009. RO-MAN 2009. The 18th IEEE International Symposium on*. IEEE, 607–612.

[11] Thibault Kruse, Amit Kumar Pandey, Rachid Alami, and Alexandra Kirsch. 2013. Human-aware robot navigation: A survey. *Robotics and Autonomous Systems* 61, 12 (2013), 1726–1743.

[12] Robert V. Levine and Ara Norenzayan. 1999. The Pace of Life in 31 Countries. *Journal of Cross-Cultural Psychology* 30, 2 (1999), 178–205. https://doi.org/10.1177/0022022199030002003 arXiv:https://doi.org/10.1177/0022022199030002003

[13] Matthias Luber, Luciano Spinello, Jens Silva, and Kai O Arras. 2012. Socially-aware robot navigation: A learning approach. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 902–907.

[14] Mehdi Moussaïd, Mubbasir Kapadia, Tyler Thrash, Robert W Sumner, Markus Gross, Dirk Helbing, and Christoph Hölscher. 2016. Crowd behaviour during high-stress evacuations in an immersive virtual environment. *Journal of The Royal Society Interface* 13, 122 (2016), 20160414.

[15] Mehdi Moussaïd, Niriaska Perozo, Simon Garnier, Dirk Helbing, and Guy Theraulaz. 2010. The walking behaviour of pedestrian social groups and its impact on crowd dynamics. *PloS one* 5, 4 (2010), e10047.

[16] Rahul Narain, Abhinav Golas, Sean Curtis, and Ming C. Lin. 2009. Aggregate dynamics for dense crowd simulation. *ACM SIGGRAPH Asia 2009 papers on - SIGGRAPH Asia 09* (2009). https://doi.org/10.1145/1661412.1618468

[17] Billy Okal and Kai O Arras. 2016. Learning socially normative robot navigation behaviors with bayesian inverse reinforcement learning. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2889–2895.

[18] Sébastien Paris, Julien Pettré, and Stéphane Donikian. 2007. Pedestrian Reactive Navigation for Crowd Simulation: a Predictive Approach. *Computer Graphics Forum* 26, 3 (2007), 665–674. https://doi.org/10.1111/j.1467-8659.2007.01090.x

[19] N Pelechano, J Allbeck, and N Badler. 2007. Controlling individual agents in high-density crowd simulation *(SCA '07)*. Eurographics Association, 99–108.

[20] Craig W Reynolds. 1987. Flocks, herds and schools: A distributed behavioral model. In *ACM SIGGRAPH computer graphics*, Vol. 21. ACM, 25–34.

[21] Mankyu Sung, Michael Gleicher, and Stephen Chenney. 2004. Scalable behaviors for crowd simulation. *Computer Graphics Forum* 23, 3 (2004), 519–528. https://doi.org/10.1111/j.1467-8659.2004.00783.x

[22] Dizan Vasquez, Billy Okal, and Kai O Arras. 2014. Inverse reinforcement learning algorithms and features for robot navigation in crowds: an experimental comparison. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 1341–1346.

[23] Mingbi Zhao, Jinghui Zhong, and Wentong Cai. 2018. A Role-Dependent Data-Driven Approach for High-Density Crowd Behavior Modeling. *ACM Transactions on Modeling and Computer Simulation* 28, 4 (2018), 1–25. https://doi.org/10.1145/3177776