# Learning to Simulate Crowds with Crowds

Bilas Talukdar, Yunhao Zhang, and Tomer Weiss
New Jersey Institute of Technology
bt26@njit.edu

## Problem

Simulating believable crowds is an open problem in graphics, with implications in movies, games and other domain such as robotics.
A crowd is composed of independent agents, where each agent selects navigational decisions that effect the present and future dynamics of the crowd.

## Related Work

- RL-based crowd simulation is an emerging area in SIGGRAPH [1, 2]
- Typical RL methods focus on manual reward tuning, which is not scalable to multiple settings.
- While Non-RL crowd simulation approaches are satisfactory for games [4], they are not able to learn to adapt to novel situations.

## Our Approach

- Our approach is learning-based, where crowd agents learn optimal navigational behavior with:

1. An RL method for learning an optimal navigational policy.
2. position-based constraints for correcting policy navigational decisions.
3. A crowd-sourcing framework for selecting policy control parameters.

## Method

### Velocity Selection

- Agents in each timestep independently select a velocity.
- Policy outputs actions as a velocity-angular velocity tuple $(v, \omega)$
- Position Based Dynamics (PBD) corrects agent's unrealistic navigational decision.

### Rewards

Our reward function consists of:

- **Distance to goal**: incentives agents to navigate towards their goals.
- **Collision Avoidance**: Penalizes overlapping agents.
- **Steering Quality**: Penalizes large changes in velocity, angular velocity, and acceleration.
- **Personal Space Maintenance**: encourages agents to keep a safe distance from others.

### Simulation Environment

- We used Position Based Dynamics (PBD) corrects agent's navigational decision [4].
- PBD adjusts agent's predicted positions when a collision is predicted.

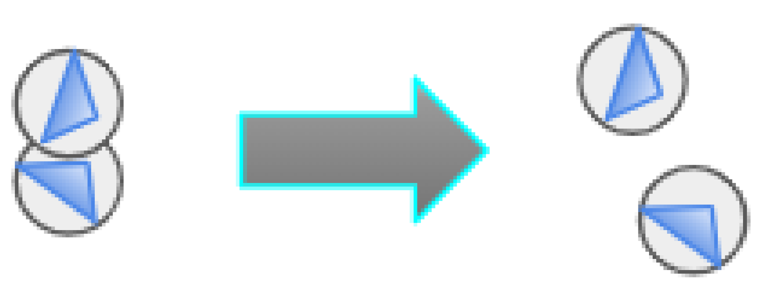

Two agents are colliding with each other, which triggers a penalty for a such encounter.

## Abstract

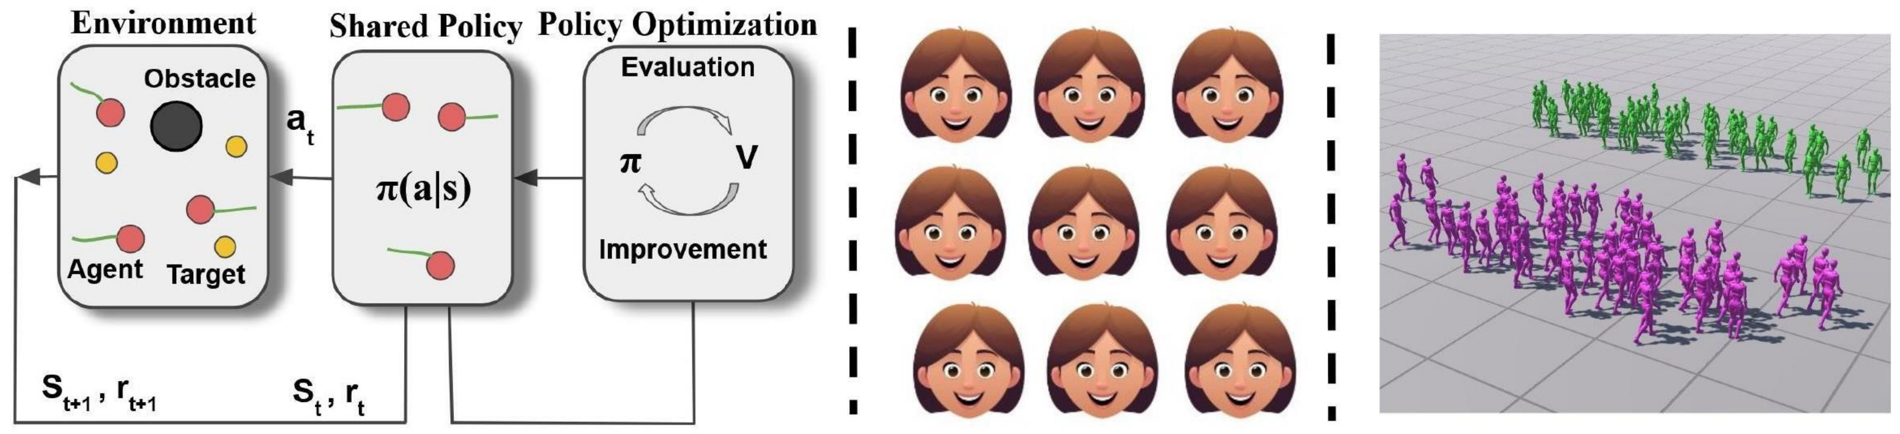


Learning  Crowd Workers  Simulation

- We propose a novel approach for simulating crowd with Deep Reinforcement Learning. Each time step t, each crowd agent selects an action $a_t$ following the policy $\pi$ and receives a reward $r_t$ from the environment. All agents share the same policy $\pi$.
- Learning optimal crowd navigation behaviors is driven by crowd-sourced feedback, which drives optimal parameters and reward for an RL policy.
- We simulate the agents with the policy trained with the optimal learning parameters suggested by the crowd-sourced Bayesian approach.

## State Representation

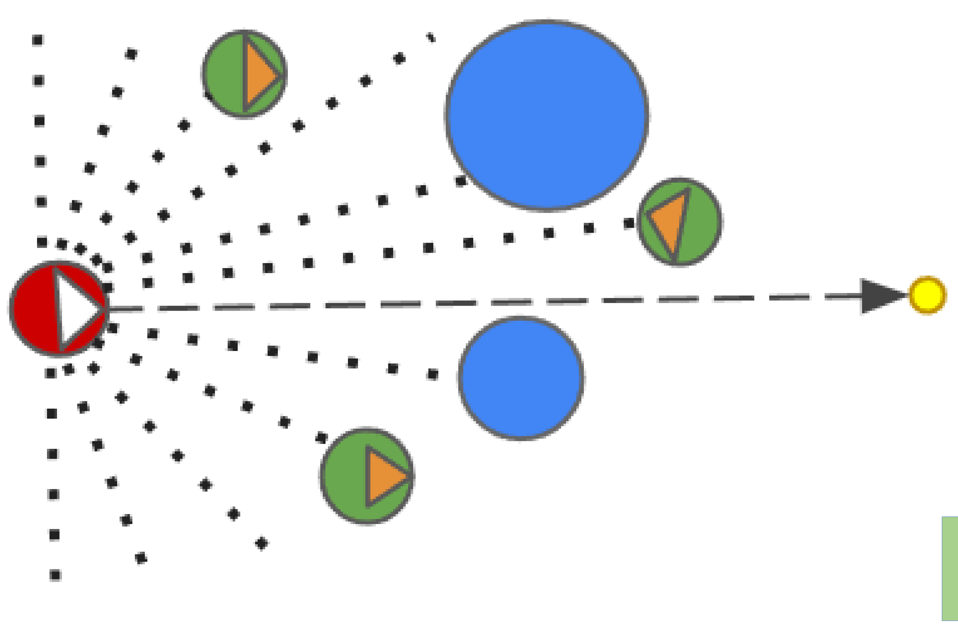

Agent sensors: an agent in red senses distance and direction to other agents in green, obstacles in blue, and its goal in yellow.

## Simulation

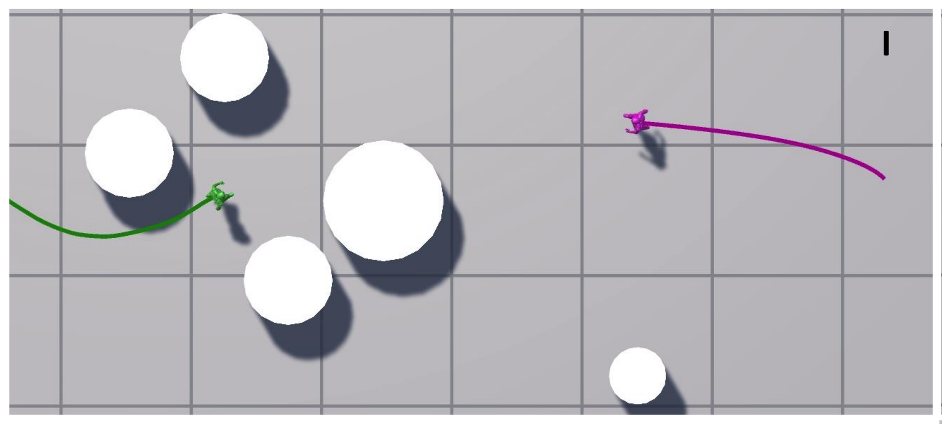


Agents locomoting in an environment containing multiple obstacles.

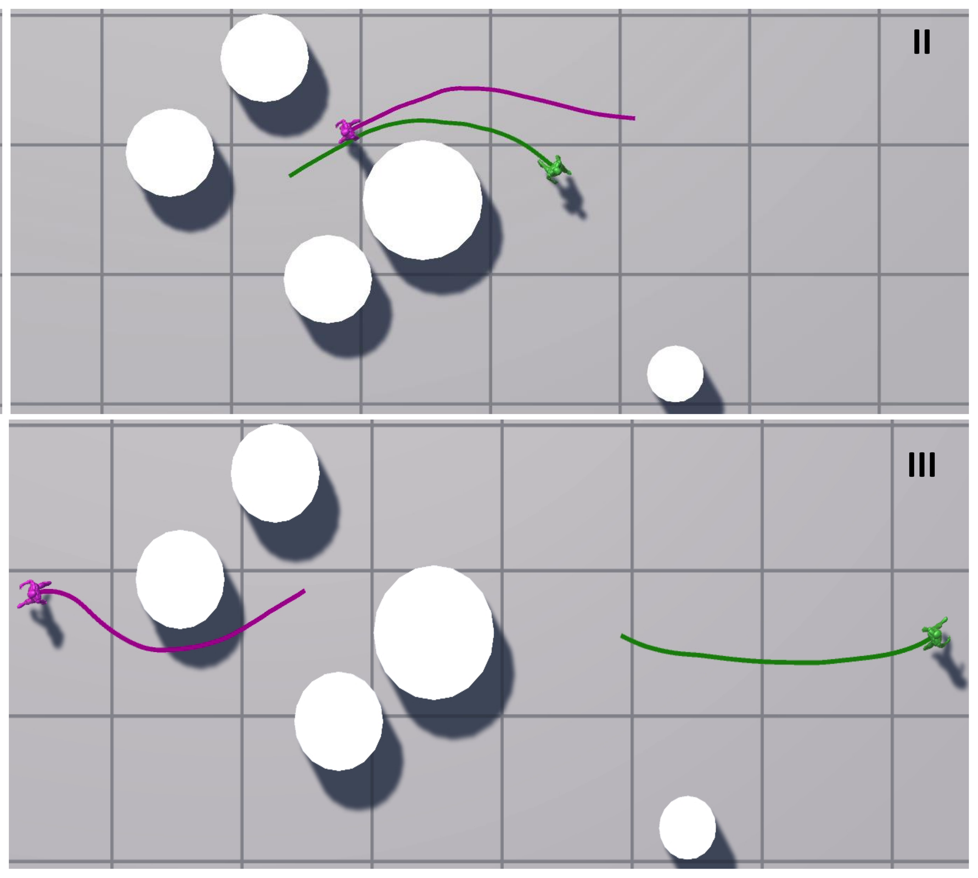


## Metrics

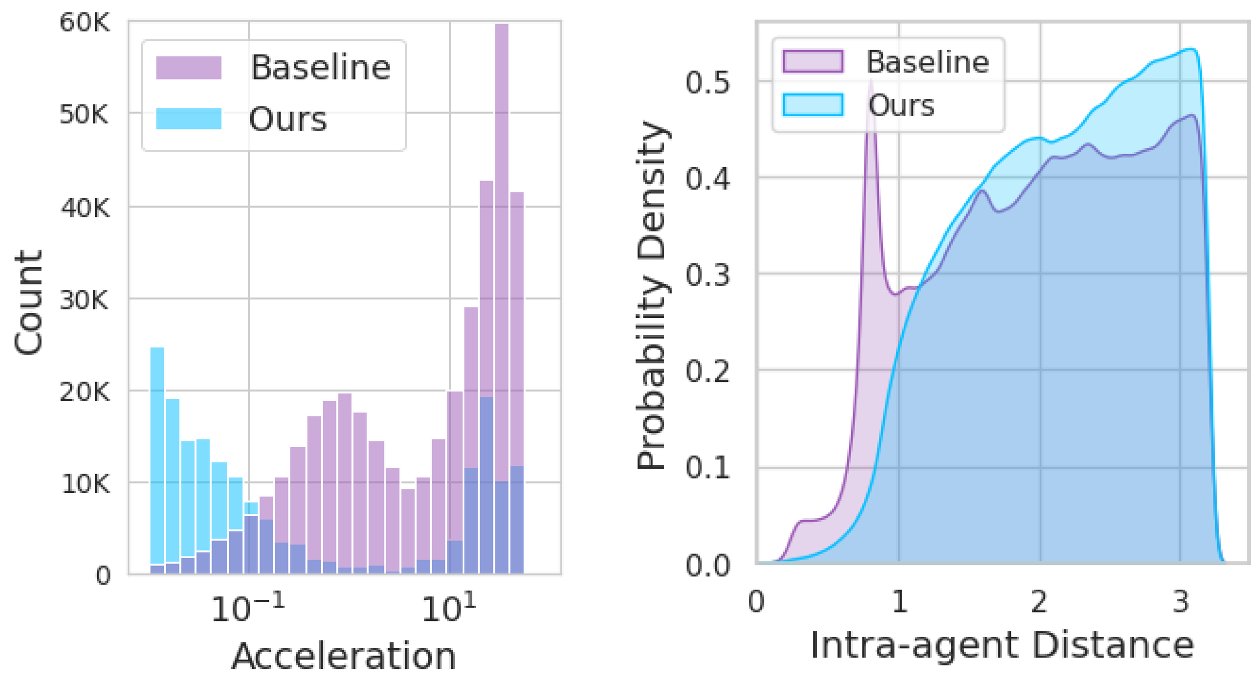


Compared to [1], agents using our method show lesser acceleration and greater inter-agent distance, which signals more realistic navigation behaviors.

## Method (Cont.)

### Learning From Preferences

- Tuning control parameters for multi-agent simulation is difficult.
- Realistic agent behavior is subjective to user perception.
- Hence, we conduct a user study to find an optimal policy.
- We used a Bayesian, Bradly-Terry model to learning preferences from comparisons [3].
- Users were given pairs of videos of multi-agent simulation and asked to select the more realistic one.
- On the preferred policy, we train a navigation policy based on the optimal control parameters inferred from user choices.

## Results And Discussion

- We trained our RL policy on a randomized hallway setting with 10 agents and 6 obstacles.
- We simulated crowds in in common benchmark scenarios (Obstacle, Circle, Hallway, etc.)
- We compared our method with i) traditional crowd simulation algorithms such as social force, and ii) recent RL proposals [1].
- Agents using our policy: i) generated smoother trajectories, ii) arrived to their destination early, iii) maintained more distance from other agents, and iv) had fewer steering changes which result in unwanted acceleration.

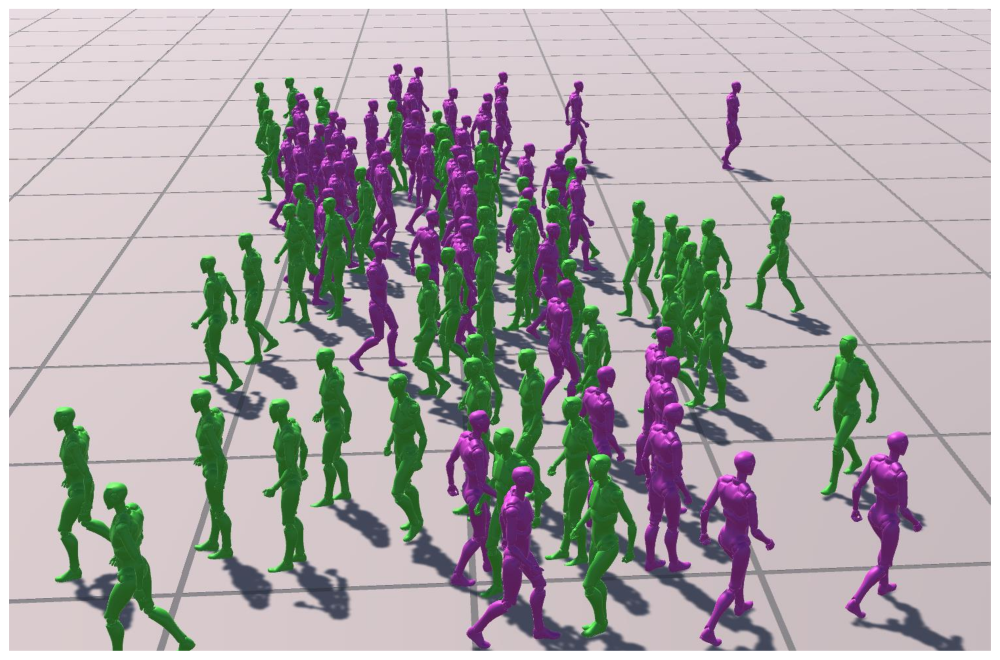

Two agent groups are passing each other using our collision avoidance approach.

## References

[1] J.Lee et al. 2018. Crowd simulation by deep reinforcement learning. In Proceedings of the 11th Annual International Conference on Motion, Interaction, and Games. 1–7.
[2] A.Panayiotou et al. 2022. CCP: Configurable Crowd Profiles. ACM Transactions on Graphics (TOG) (2022).
[3] Y. Koyama et al. 2017. Sequential line search for efficient visual design optimization by crowds. ACM Transactions on Graphics (TOG) 36 (2017), 1 – 11
[4] T.Weiss et al. 2019 "Position-based real-time simulation of large crowds." Computers & Graphics 78: 12-22.