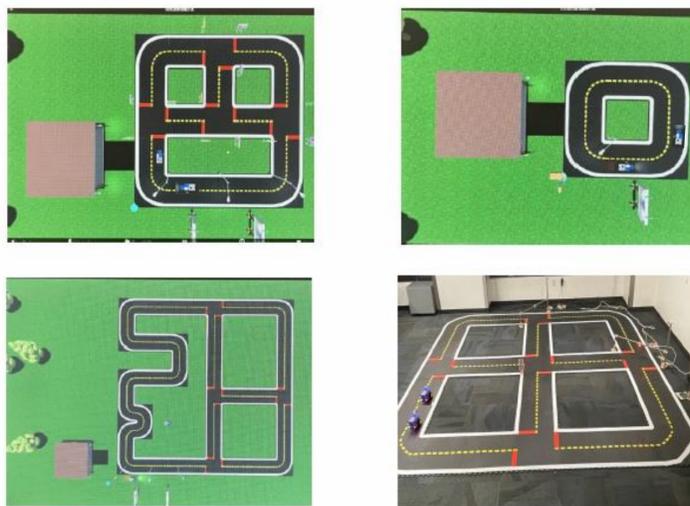


Introduction

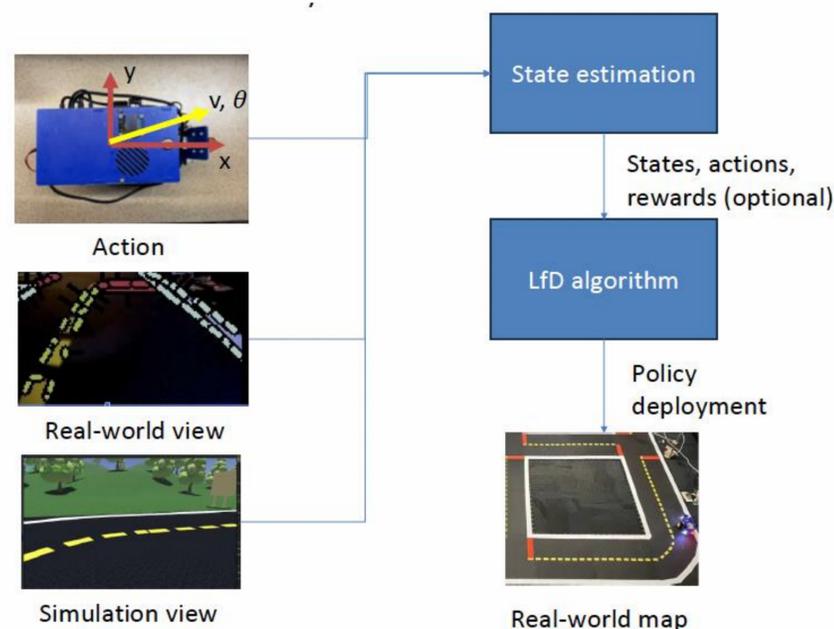
- **Problem:** The sim-to-real problem, referring to the difficulty of transferring policies learned in simulated environments to real-world scenarios, is a critical challenge in autonomous systems and robot learning. While simulations provide safe, low-cost training, the reality gap limits policy transfer to real-world settings.
- **Method:** We hypothesize that **Learning from Demonstration** (LfD) and related methods using task-relevant external data will facilitate effective knowledge transfer from simulation to real-world environments.
- **Impact:** Our work investigates which RL algorithms are effective, impact of using simulation and real-world data, and establishes a baseline for autonomous driving tasks



Simulator and real-world maps

Experiment Design

- **Task:** The robot is required to perform lane switching when approaching a vehicle
- **Data collection:** We collected training data from three sources: simulation-only (S), real-world-only (R), and a mixture of both (S+R). The expert demonstrations included scenarios with no obstacles as well as situations where a vehicle is present in front of the robot.
- **Methodology:** We used imitation learning, specifically behavior cloning to train the policy.
- **Evaluation Metrics:** We evaluated performance using survival time and success rate across three trials of an obstacle avoidance task. A trial ends when the robot deviates from the lane.



Current progress

	Average Success Rate (%)	Average Survival time (s)
S	33	5
R	100	9
S+R	0	5

Observation and Future work

- **Observation:** Currently, only real-world data appears to improve the robot's performance. This suggests that the performance gap may be due to differences between the simulation and real-world environments, or limitations in the imitation learning approach when trained on simulated data and deploy in real-world scenario.
- **Future works:** We plan to explore sim-to-real techniques such as domain randomization and domain adaptation. Additionally, we will investigate hyperparameter tuning and advanced reinforcement learning methods, including inverse reinforcement learning, to improve both performance and generalization.

The work was partially supported by NSF awards IIS-2421839, the OpenAI Researcher Access program, and Amazon Research Award. The views and conclusions contained in this paper are those of the authors and should not be interpreted as representing any funding agencies.