ISAAC SHEIDLOWER

177 College Ave, Medford, MA, 02155 | Website https://isheidlower.github.io | Github: https://github.com/IsaacSheidlower | Isaac.Sheidlower@Tufts.edu

EDUCATION

Tufts University
PhD: Department of Computer Science **2020-Current (Expected graduation date: Spring 2025)**
Advisor: Dr. Elaine Short
Relevant course work: Software Engineering; Statistical Pattern Recognition; Human-Robot Interaction; Probabilistic Robotics

Rutgers University New Brunswick
Bachelor of the Arts **2016-2020**
Major: Cognitive Science; Minor: Computer Science
Relevant Couse Work: Artificial Intelligence; Language as Data; Intro to Neuromorphic Computing
Honors: Rutgers New Brunswick Honors Program, Phi Beta Kappa

PUBLICATIONS AND PAPERS

I. Sheidlower, Emma Bethel, Douglas Lilly, Reuben M. Aronson, and Elaine Schaertl Short. "Imagining In-distribution States: How Predictable Robot Behavior Can Enable User Control Over Learned Policies" (Accepted for publication at RO-MAN 2024) https://arxiv.org/pdf/2406.13711

I. Sheidlower, Reuben M. Aronson, and Elaine Schaertl Short. "Towards Interpretable Foundation Models of Robot Behavior: A Task Specific Policy Generation Approach" (Workshop paper at RLC 2024, Training Agents with Foundation Models)

I. Sheidlower, Mavis Murdock, Emma Bethel, Reuben M. Aronson, and Elaine Schaertl Short. "Online Behavior Modification for Expressive User Control of RL-Trained Robots." In *Proceedings of the 2024 ACM/IEEE International Conference on Human-Robot Interaction*, 639–48. HRI '24. New York, NY, USA: Association for Computing Machinery, 2024. https://doi.org/10.1145/3610977.3634947.

I. Sheidlower, A. Moore and E. Short, *"Keeping Humans in the Loop: Teaching via Feedback in Continuous Action Space Environments," 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Kyoto, Japan, 2022, pp. 863-870, doi: 10.1109/IROS47612.2022.9982282. https://ieeexplore.ieee.org/document/9982282*

I. Sheidlower, A. Moore, and E. Short. "Environment Guided Interactive Reinforcement Learning: Learning from Binary Feedback in High-Dimensional Robot Task Environments." In Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems, 1726–28. AAMAS '22. (Extended Abstract)

I.S. Sheidlower and E. Short. "When Oracles Go Wrong: Using Preferences as a Means to Explore". Companion of the ACM/IEEE Conference on Human-Robot Interaction (HRI): Late-Breaking Reports, Online,
2021. **Best LBR Nominee.**

RESEARCH INTERESTS

My research centers around empowering users through state-of-the-art robot learning algorithms. My approach to this entails human-in-the-loop teaching, explicitly considering the end user when training a robot offline, and making use of simulation to equip robots with robot policies that people want to use and experiment with.

RESEARCH PROJECTS AND EXPERIENCE

Assistive Agent and Behavior Lab at Tufts University

- *Enabling Novel Task Execution with User Control and Imagined Robot States (2023-)*
    - *Summary*: Given an RL policy and a means of teleoperation, as a user simultaneously teleoperates the robot and an autonomous policy executes, a user may bring the **policy out-of-distribution** or lead to failure with regard to the robot's reward function. This can lead to unpredictable behavior. We propose using out-of-distribution detection and simulation rollouts to resolve potentially unpredictable robot behavior in response to user control.
- *Customizing the Behavior of an RL based robot* (2022-2023)
    - *Skills*: Python, Pytorch, ROS, OpenCV, NumPy, Mujoco, Bullet, Gazebo, OpenAI Gym, Brax Gym, Kinova Robot Arms (Jaco Gen2 and Gen3)
    - *Summary*: While optimal policies may successfully complete a task, they may not meet a user's needs for how that task should be completed. For example, a user may wish to control the robot's speed of movement in a locomotion task or

style of painting in a painting task. We developed a novel diversity-inspired algorithm, Adjustable Control Of Rl Dynamics (ACORD), which grants users real-time control over aspects of a robot's behavior as it executes an RL policy. We designed a sim-to-real pipeline for the Kinova Gen3 arm using ROS and Gazebo for a painting task and deployed that task in a user study where we compared different ways a user can control a robot as it completes a painting task. We found ACORD has preferred levels of expression over pure RL with similar or greater levels of control as Shared Autonomy.

- *Interactive Reinforcement Learning in Continuous Action Spaces* (2020-2022)
  - *Skills*: Python, Pytorch, Flask, HTML, JavaScript, OpenCV, NumPy, Scipy, Mujoco, Bullet, Gazebo, OpenAI Gym, Brax Gym, AWS, Amazon Mechanical Turk
  - *Summary*: Interactive Reinforcement Learning (IntRL/RLHF) has been a powerful way to teach robot's new tasks just through binary feedback. However, there is a lack of IntRL algorithms that operate in continuous action spaces and are competitive with state-of-the-art RL. To address this, we introduced Continuous Action-space Interactive Reinforcement Learning (CAIR), an Interactive Reinforcement Learning that can learn from teachers providing good/bad feedback to a robot in high-dimensional continuous action spaces. We ran an online user-study through MTurk where we collected human feedback in real time as a simulated robot learned to perform a task. This involved deploying an off-policy RL algorithm on a server and deploying a real-time robotics simulation for users to interact with. ACORD outperformed both standard RL and prior IntRL approaches.
  - *Resulting publications*: "Keeping Humans in the Loop: Teaching via Feedback in Continuous Action Space Environments" (IROS 2022); "Environment Guided Interactive Reinforcement Learning: Learning from Binary Feedback in High-Dimensional Robot Task Environments." (AAMAS 2022, extended abstract).
- *Learning From Teachers with Different Preferences* (2020)
  - *Skills*: [Python, NumPy, Scipy, Matplotlib]
  - *Summary*: Prior algorithms in Interactive Reinforcement Learning Proposed suffered when multiple teachers with different preferences would teach the same agent. We proposed interactive reinforcement learning algorithm that could leverage differences in preferences among teachers to improve task robustness and learning. We also proposed a theoretical model of how teachers give binary evaluative feedback to robots
  - *Resulting publications*: "When Oracles Go Wrong: Using Preferences as a Means to Explore." (HRI 2021, late-breaking reports).
- *Learning From Teachers with Different Levels of Authority* (2020, Tufts Human-robot Interaction class)
  - *Skills*: MicroPython, NumPy, Scipy, Matplotlib, EV3 LEGO robot
  - *Summary*: Proposed and tested an algorithm that incorporates information about a user's hierarchical status in the context of the learning robot (e.g., feedback from a parent and child in a household setting, or a caregiver and a patient)
  - *Evaluation*: Deployed the algorithm on LEGO EV3 and successfully learned a task by prioritizing feedback from a user with more authority over feedback from a user with less authority

REVIEWER

Conferences: Int. Conference on Autonomous Agents and Multiagent Systems (AAMAS), ACM/IEEE Conference on Human-Robot Interaction (HRI)

TEACHING

Tufts Course Teaching:

Introduction to ROS – Spring 2023

Tufts Teaching Assistant:

Comp 150/ME 193: Assistive Algorithms - Spring 2021

MENTORING

Tufts Undergraduate Students:

Emma Bethel (January 2023-), working on developing novel assistive robot and novelty detection algorithms
Alina Shah & Teo Patrosio (May 2022-), working on developing a codebase for a Fetch Freight100
Allison Moore (Sep. 2021 – May 2022), worked on data visualization techniques for educational and research purposes

RELATED EXPERIENCE

ACM/IEEE Conference on Human-Robot Interaction (HRI)

Student Volunteer **2021**

Assisted in editing camera-ready PDFs to be more accessible.