

Reinforcement Learning for Robust Robot Skills

Faculty advisor: Prof. Chi-Guhn Lee

Reinforcement learning has the potential to produce robot behaviours that are adaptive, robust, and difficult to hand-engineer. However, applying RL to real robotic systems remains challenging due to sample efficiency, safety, sim-to-real gaps, reward design, and reliable evaluation.

This project explores reinforcement learning for robotic manipulation and control, with a focus on developing and evaluating robot skills in simulation and transferring insights toward real-world systems. Students may work on reward design, policy learning, curriculum learning, sim-to-real transfer, benchmarking, and comparisons with imitation learning or other robot learning approaches. The project may involve experiments using the Trossen Stationary AI Robot, UR5 robot, and other real robotic platforms.

Preferred Foundation / Experience

- Strong foundation in reinforcement learning and/or robot learning
- Proficiency in Python and PyTorch
- Familiarity with simulation environments such as IsaacSim, IsaacLab, MuJoCo, or Gym-style frameworks
- Interest in manipulation, control, evaluation, and sim-to-real deployment
- Comfort running experiments, analyzing results, and tuning learning systems
- Prior coursework or projects in RL, robotics, or deep learning is a strong plus

Why this project?

You'll work on a technically challenging and highly relevant area of robot learning, gaining practical experience with RL pipelines, simulation, and real robotic systems. This project is especially suitable for students aiming for robotics ML roles, applied AI research, or future graduate work.

Contact: Ramy ElMallah, ramy.elmallah@mail.utoronto.ca; Chongyu Zhu, chongyu.zhu@mail.utoronto.ca