

Faculty of Engineering

Summer Research Program 2024-2025

Project Title:

Reinforcement Learning with Opponent Modeling in Repeated Rock-Paper-Scissors

Supervisor(s): Chenyuan Zhang, Dana Kulic, Michael Burke, Pamela Carreno-Medrano

Department: Electrical and Computer Systems Engineering

Email: chenyuan.zhang@monash.edu

Website profile of project supervisor: <https://sites.google.com/view/chenyuanzhang>

Objective

The study aims to investigate how understanding and predicting opponents' behavior can enhance the learning agent's ability to adapt and win over successive rounds, compared to traditional reinforcement learning approaches that do not account for opponent behavior.

Project Details

In two-player competitive games, understanding and predicting opponents' behavior is crucial for winning. Ideally, reinforcement learning algorithms can implicitly encode opponent information in the policy by including the playing history, but this requires a substantial amount of data and training time to converge. However, evidence shows that humans can identify their opponents' strategies in a more efficient manner in Repeated Rocker-Paper-Scissors (RRPS) ([Brockbank & Vul. 2021](#)). Inspired by this, [Lanctot et al. \(2023\)](#) observed that an explicit opponent identification step improves the performance of the RL agent on RRPS. However, their work assumes a fixed pool of potential opponents, which is not typically available in real-world settings.

In this project, we aim to extend Lanctot et al.'s work by replacing the supervised opponent modeling with an unsupervised algorithm, which can adaptively generate a low-dimensional policy representation of the opponent based on the playing history ([Grover et al. 2018](#)) and use that representation to predict the next move. This approach does not require a pre-defined opponent population and offers better generalization to other domains.

The project will begin by replicating the work of [Lanctot et al. \(2023\)](#) using the RL framework [OpenSpiel](#), and then integrate policy representation learning, testing it against the same set of bot populations to assess its effectiveness.

Prerequisites

Experience with Reinforcement Learning; familiar with Python Programming

Additional Information

Applicants may be required to attend an interview