Reinforcement Learning for RoboCup Soccer Keep Away

The RoboCup Small Size League (SSL) is one of the oldest RoboCup Soccer leagues. It is focused on intelligent multi-agent coordination and control in a highly dynamic environment. Teams currently compete in 11 vs. 11 soccer matches (6 vs. 6 in division B) with an orange golf ball. The robots are built by the teams and have to conform to certain dimensions and constraints, as defined in the rules. All objects on the field are tracked by a central vision system – called SSL-Vision – which is maintained by the SSL community. Off-field computers for each team are used for the processing required for coordination and control of the robots. Communication is wireless and uses dedicated commercial radio transmitter/receiver units. (Source: https://ssl.robocup.org/)

Unlike other RoboCup leagues, the Small Size League (SSL) has seen limited application of reinforcement learning (RL) techniques. This project aims to explore the potential benefits of utilizing RL methods in the SSL by focusing on a specific task: a "keepaway" scenario. We propose to demonstrate how state-of-the-art RL algorithms can be effectively employed to enhance team performance in dynamic, real-time environments.

The proposed "keepaway" task is inspired by the "Reinforcement Learning for 3 vs. 2 Keepaway" paper (https://link.springer.com/chapter/10.1007/3-540-45324-5_23). In this scenario, one team (forwards) is tasked with maintaining possession of the ball within a predefined area, while the opposing team (defenders) attempts to gain possession. The episode ends when the defenders successfully gain control of the ball or when the ball exits the region, after which the players are reset for a new episode.

The objective of this project is to apply advanced reinforcement learning techniques to train agents in the "keepaway" task, aiming to improve their ability to maintain possession and adapt to the defenders' strategies. By leveraging RL, we expect to reveal insights into how AI can improve tactical decision-making and performance in the SSL.

Requirements

- Solid experience in Reinforcement Learning (RL)
- Preferably experienced with physics simulation environments (e.g., Isaac Gym)
- Strong interest in RoboCup Soccer and robotics
- Proficiency in PyTorch

This project is part of a collaboration with an Australian team called TurtleRabbit. You can have a look at their team description paper for the RoboCup 2024 (https://arxiv.org/pdf/2402.08205) and at their website (https://www.turtlerabbit.org/).

If you are interested in this project feel free to contact me!

Contact
Lisa Graf (lgraf@informatik.uni-freiburg.de)