Motion Planning for Structured Exploration in Robotic Reinforcement Learning

William Cannon Lewis II, Rice University

Reinforcement Learning

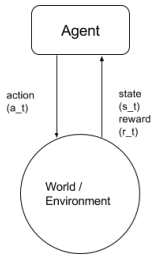

In Reinforcement Learning (RL), we are interested in learning a **policy** which maximizes **reward** received from an **environment**. In robotic RL, that environment is a robot interacting with either a simulation or neural network to parameterize any part of an RL algorithm, it is called **Deep RL**. Deep RL methods are particularly good at dealing with the high-dimensional, continuous state and action spaces typically encountered in robotic RL.

Exploration

A common problem in RL is the **exploration-exploitation** tradeoff. Typically, exploration is done by either injecting noise into the policy that an agent follows, or by selecting a random action periodically with a certain probability.

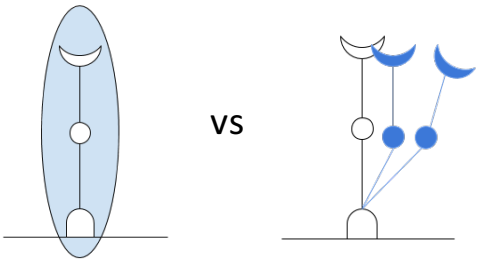

However, this strategy doesn't work well for robots, since robots have such complex configuration spaces. Robotic tasks are also typically more complex than classical RL tasks, and so need more directed exploration. By using motion planning, we can gain experience from long, structured trajectories and hence sequences of states, actions, and rewards.

OpenAI Gym, ROS, and OMPL

To do this research, we needed a robot simulation environment, a way to do motion planning, and an interface for reinforcement learning with general algorithms.

OpenAI Gym BETA

ROS

OMPL

- **OpenAI Gym** is a standard library for RL which allows for defining new environments.
- **ROS** is a popular suite of tools for working with robotics, including the Gazebo simulator.
- **OMPL** is the Open Motion Planning Library, built and maintained by Rice.

These tools allow us to test many different ways of using motion planning for exploration in robotic RL, though the work is still in progress at this point.

Current Work

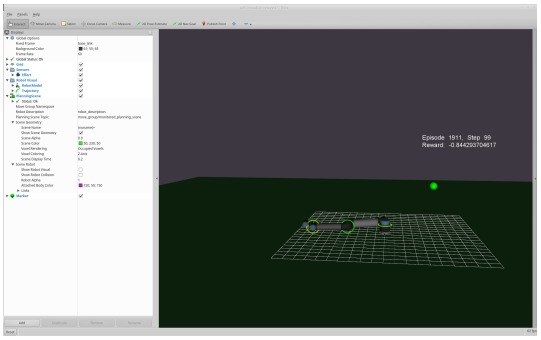

If you are interested in using any of these tools or getting involved in this project, please let me know at wcl2@rice.edu!