Motion Planning for Structured Exploration in Robotic Reinforcement Learning

Cannon Lewis (wcl2@rice.edu), Rice University

Reinforcement Learning

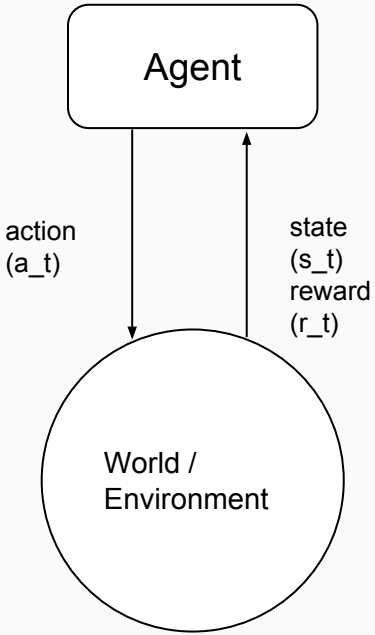

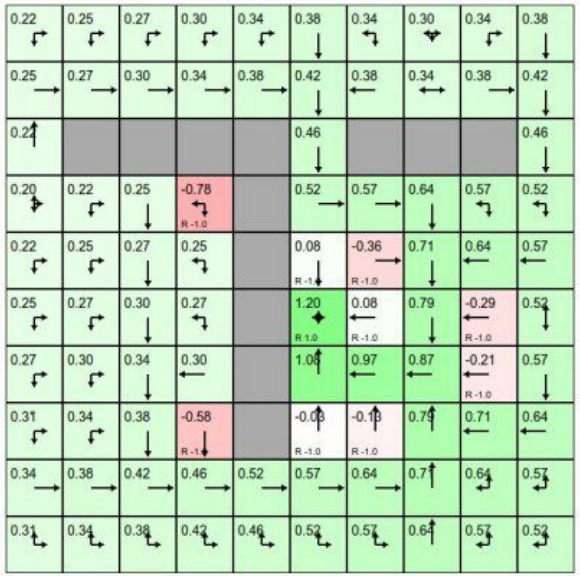

http://cs.stanford.edu/people/karpathy/reinforcejs/index.html

Deep Reinforcement Learning

Algorithms (typically value-function based) which use a neural network as a value function or policy approximator.

- Examples are DQN, TRPO, DDPG, A3C, etc.
- Much better than traditional RL approaches at handling continuous state and action spaces.
- Responsible for most recent high-profile advances in RL.

ROS Gym

Goal: Make ROS accessible to agents using the OpenAI Gym API

- Actions and States are ROS nodes
  - Can run learning in the Gazebo simulator or on a real robot (not very safe)
- Two parts:
  - Robot Configuration:
    - Joints to control (action space)
    - Sensors + joint states (state space)
  - Reward Function: depends on the task you want the robot to accomplish

Ros Gym Demonstration

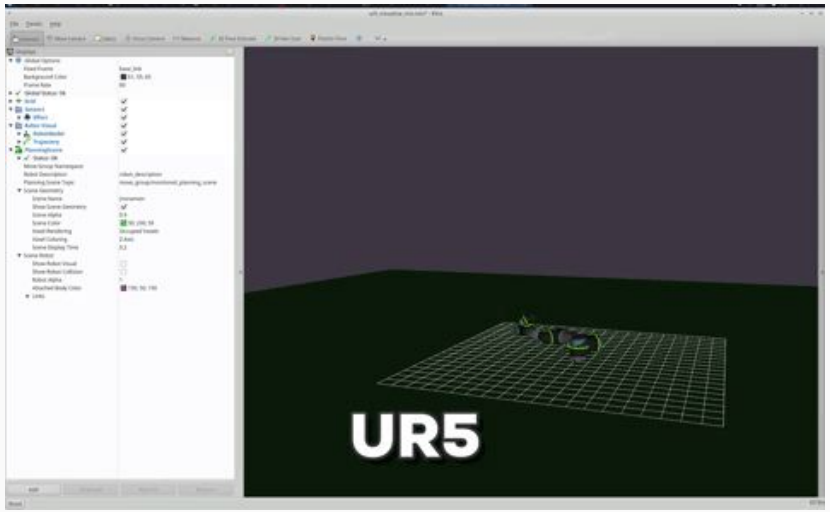

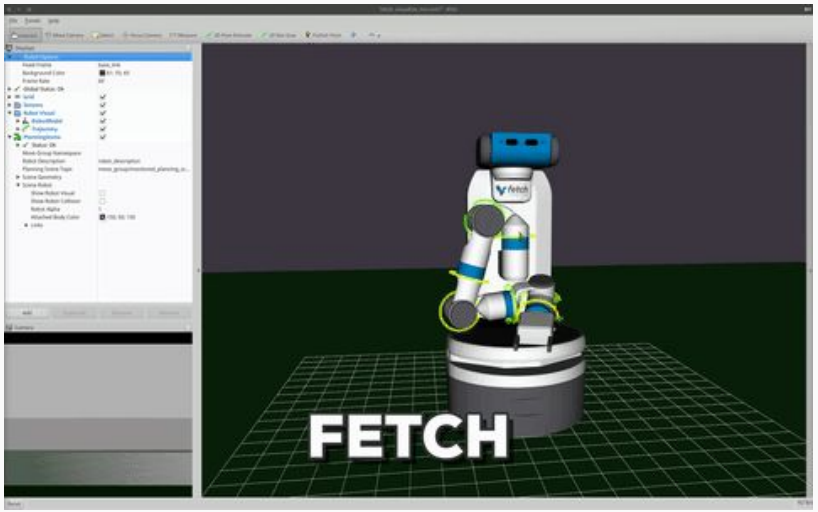

Robotic RL is Hard

- Traditional RL problems are 2D or 3D and discrete.
- Robots have high dimensional (~6D - 18D) and continuous state and action spaces.
- For previous work in robotic RL, researchers have relied on discretization or dynamical motor primitives.
    - However, these simplifications are undesirable as they limit the applicability of learned robot policies.

Motion Planning for RL

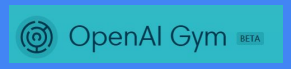

- One of the biggest problems for robotic RL is exploration.
  - Typically, exploration is done by taking random actions.
  - With complex robot dynamics, this exploration is ineffective.
- Motion Planning allows us to construct longer, directed trajectories for exploration.
- Using MoveIt and OMPL, we enable RL algorithms to do structured exploration on robots.

Motion Planning for Robotic RL

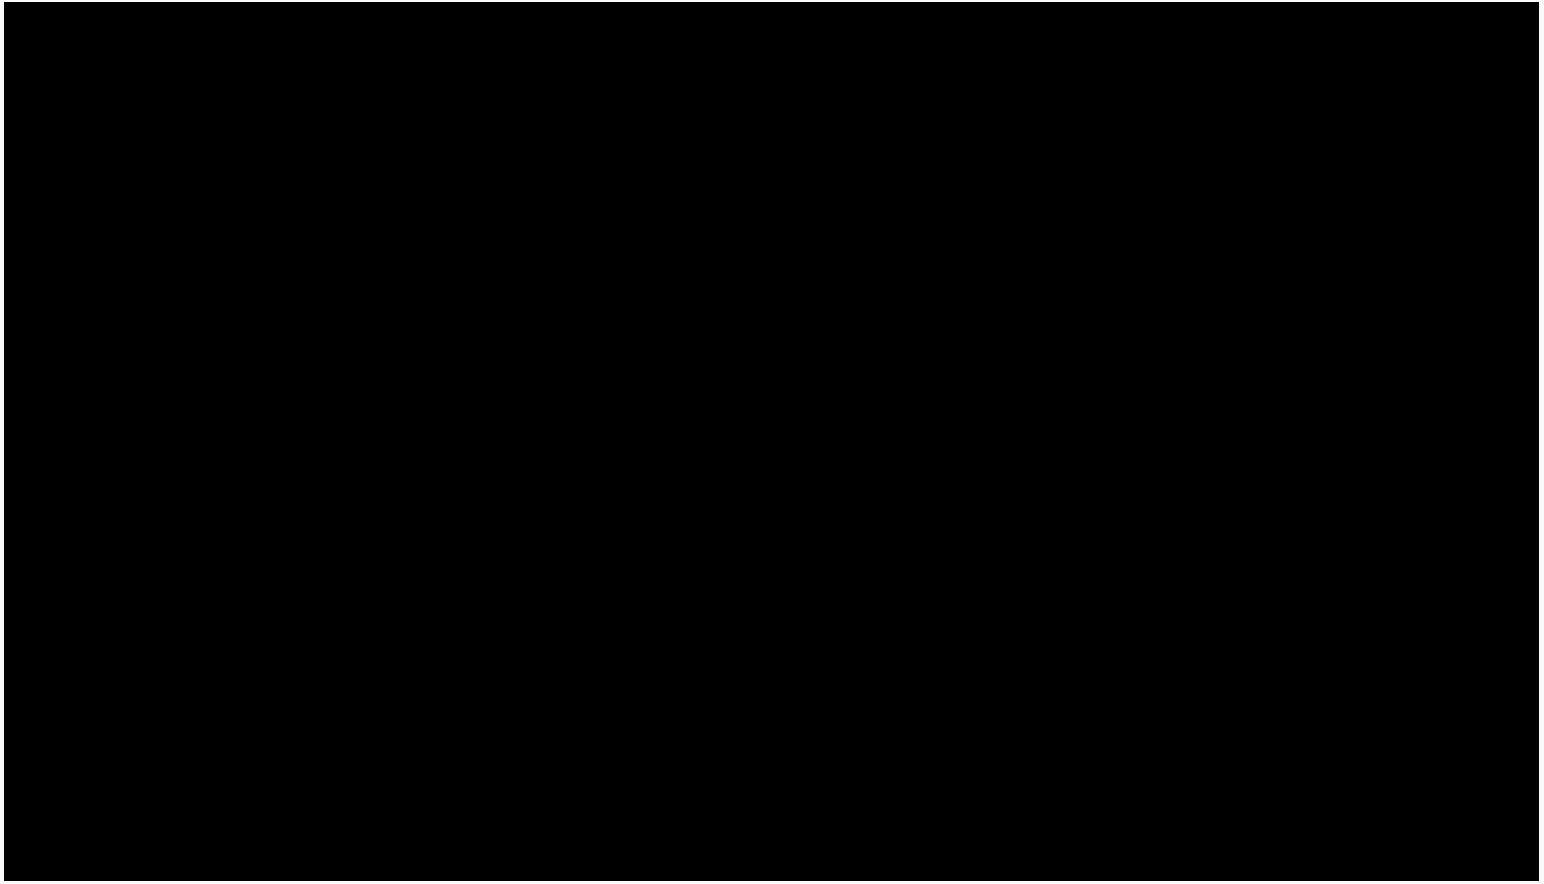

Thank You