

Homework 10

Using the assigned reading listed on the course page, answer the questions below with a short response. Note that we are looking for concise statements that show understanding, not quantity. The total discussion should roughly be a page.

Deep Reinforcement Learning and Robotics

1. Reinforcement learning (RL) is generally considered to be harder than supervised learning. What are the fundamental differences from supervised learning that make RL hard? (Feel free to use any definition of RL – learning from rewards, learning via interaction with an environment, $\arg \max_{\pi} \mathbb{E}_{\tau \sim \pi} [\sum_t \gamma^t r(s_t, a_t)]$, etc – just state your assumptions).
2. What is the Q function estimating? Why is this a useful quantity to estimate and store? In what way does the Q function save computation?
3. Ibarz et al. discuss both model-free and model-based approaches to deep reinforcement learning for robotics. What are some of the trade-offs they identify between these two approaches?
4. Pick a robotic task you find interesting. For example, you could pick the task of constructing a piece of Ikea furniture, or the task of driving to a New England ski resort in the winter, or the task of searching for a lost person in the woods, or a task you make up yourself. Propose a general plan for how you would tackle this problem using deep RL. Refer to concepts from the readings and lectures – for example, would you pick a model-free or a model-based approach, would you train in simulation or on the physical robot itself, etc.

Submission: Upload a PDF of your response through Canvas by **11/23 at 1pm**.