

THESIS DEFENSE

Robot Deep Reinforcement Learning: Tensor State-Action Spaces and Auxiliary Task Learning with Multiple State Representation

Tuesday, August 4, 2020

10:00 a.m. ET

Thesis Committee:



Manuela Veloso
Chair



Katerina Fragkiadaki



David Held



Martin Riedmiller
Google DeepMind



Devin Schwab

Abstract

A long standing goal of robotics research is to create algorithms that can automatically learn complex control strategies from scratch. Part of the challenge of applying such algorithms to robots is the choice of representation. Reinforcement Learning (RL) algorithms have been successfully applied to many different robotic tasks such as the Ball-in-a-Cup task with a robot arm and various RoboCup robot soccer inspired domains. However, RL algorithms still suffer from issues of large training time and large

amounts of required training data. Choosing appropriate representations for the state space, action space and policy can go a long way towards reducing the required training time and required training data.

This thesis focuses on robot deep reinforcement learning. Specifically, how choices of representation for state spaces, action spaces, and policies can reduce training time and sample complexity for robot learning tasks. In particular the focus is on two main areas:

1. Transferrable Representations via Tensor State-Action Spaces
2. Auxiliary Task Learning with Multiple State Representations

The first area explores methods for improving transfer of robot policies across environment changes. Learning a policy can be expensive, but if the policy can be transferred and reused across similar environments, the training costs can be amortized. Transfer learning is a well-studied area with multiple techniques. In this thesis we focus on designing a representation that makes for easy transfer. Our method maps state-spaces and action spaces to multi-dimensional tensors designed to remain a fixed dimension as the number of robots and other objects in an environment varies. We also present the Fully Convolutional Q-Network (FCQN) policy representation, a specialized network architecture that combined with the tensor representation allows for zero-shot transfer across environment sizes. We demonstrate such an approach on simulated single and multi-agent tasks inspired by RoboCup Small Size League (SSL) and a modified version of Atari Breakout. We also show that it is possible to use such a representation and simulation trained policies with real-world sensor data and robots.

The second area examines how strengths in one robot Deep RL state representation can make-up for weaknesses in another. For example, we would often like to learn tasks using the robot's available sensors, which include high-dimensional sensors such as cameras. Recent Deep RL algorithms can learn with images, but the amount of data can be prohibitive for real robots. Alternatively, one can create a state using a minimal set of features necessary for task completion. This has the advantages of 1) reducing the number of policy parameters and 2) removing irrelevant information. However, extracting these features often has a significant cost in terms of engineering, additional hardware, calibration and fragility outside the lab. We demonstrate this on multiple robot platforms and tasks in both simulation and the real-world. We show that it works on simulated RoboCup Small Size League (SSL) robots. We also demonstrate that such techniques allow for from scratch learning on real hardware via the Ball-in-a-Cup task performed by a robot arm.