



Electrical Engineering and Computer Science

Ph.D. Dissertation Defense

Learning Representations in Reinforcement Learning

Jacob Rafati

Electrical Engineering and Computer Science
University of California, Merced

Abstract:

Reinforcement Learning (RL) algorithms allow artificial agents to improve their action selection policy to increase rewarding experiences in their environments. Temporal Difference (TD) learning algorithm, a model-free RL method, attempts to find an optimal policy through learning the values of agent's actions at any state by computing the expected future rewards without having access to a model of the environment. TD algorithms have been very successful on a broad range of control tasks, but learning can become intractably slow as the state space grows. This has motivated methods for using parameterized function approximation for the value function and developing methods for learning internal representations of the agent's state, to effectively reduce the size of state space and restructure state representations in order to support generalization. This dissertation investigates biologically inspired techniques for learning useful state representations in RL, as well as optimization methods for improving learning. There are three parts to this investigation. First, failures of deep RL algorithms to solve some relatively simple control problems are explored. Taking inspiration from the sparse codes produced by lateral inhibition in the brain, this dissertation offers a method for learning sparse state representations. Second, the challenges of RL in efficient exploration of environments with sparse delayed reward feedback, as well as the scalability issues in large-scale applications are addressed. The hierarchical structure of motor control in the brain prompts the consideration of approaches to learning action selection policies at multiple levels of temporal abstraction. That is learning to select subgoals separately from action selection policies that achieve those subgoals. This dissertation offers a novel model-free Hierarchical Reinforcement Learning framework, including approaches to automatic subgoal discovery based on unsupervised learning over memories of past experiences. Third, more complex optimization methods than those typically used in deep learning, and deep RL are explored, focusing on improving learning while avoiding the need to fine tune many hyperparameters. This dissertation offers limited-memory quasi-Newton optimization methods to efficiently solve highly nonlinear and nonconvex optimization problems for deep learning and deep RL applications. Together, these three contributions provide a foundation for scaling RL to more complex control problems through the learning of improved internal representations.

Biography:

Jacob Rafati is a Ph.D. candidate in the EECS program at the UC Merced. He is also a member of Dr. David C. Noelle's Computational Cognitive Neuroscience Laboratory. His research focus is on Large-scale Numerical Optimization, Machine Learning and Reinforcement Learning. In his Ph.D. dissertation, he has investigated, and implemented biologically inspired methods for learning representations in model-free reinforcement learning framework. He has also studied optimization methods for large-scale machine learning problems, such as deep learning and deep reinforcement learning. For more details, and publications visit his website at <http://rafati.net>.

Date:

04 / 26 / 2019

Time:

1:00 PM to
2:00 PM

Location:

Social Science and
Management - 317
Half Dome Conference Room
(3rd floor of SSM BLDG)

More Information:

Jacob Rafati
<http://rafati.com/>
jrafatiheravi@ucmerced.edu

Faculty Advisor:

David C. Noelle
dnoelle@ucmerced.edu