

**Title:** Deep Reinforcement Learning

**Abstract:** Autonomous agents has been recently widely used to solve complex problems and used for navigation in complex environments. In this tutorial we introduce the main concepts behind the Reinforcement Learning (RL) and how to define Markov Decision Process. We will talk about the state-of-the-art algorithms and approaches that are currently used to solve RL problems. And also introducing practical examples of the main algorithms.

**Agenda:**

**Section 1: Fundamentals of Reinforcement learning [80 minutes + 10 minutes break]**

- Introduction and core elements of RL
- Recent advances and applications of RL
- Defining RL framework and Markov Decision Process
- Core algorithms in RL (SARSA, Q-learning, DQN)

**Section 2: Deep RL and Practical Applications [60 minutes]**

- Policy gradient approximation algorithms (TRPO, PPO)
- Actor-critic methods (A3C, A2C)
- Coding examples of RL algorithms

**Expected audience:** We expect two types of audience. First type is attendees who have backgrounds in computer science, engineering, and other mathematically oriented disciplines. Second type is interested behavioral and social scientists, who will benefit by understanding the main reinforcement learning concepts. We will go over some code examples, so laptop is recommended, but not required. The tutorial will be based on the graduate-level courses "Introduction to Reinforcement Learning" and "Stochastic Simulation and Inference" taught by Alina Vereshchaka and Dr. Wen Dong accordingly.

**Speakers bio:**

Alina Vereshchaka is a PhD student of Computer Science and Engineering at the State University of New York at Buffalo, USA. Her current research interests include deep reinforcement learning, optimization and multi-agent modeling in stochastic environments. She developed and leading a course for graduate students in Reinforcement Learning and undergraduate course in Artificial Intelligence during Summer 2019 at the State University of New York at Buffalo, USA. Email address is [avereshc@buffalo.edu](mailto:avereshc@buffalo.edu)

Dr. Wen Dong is an Assistant Professor of Computer Science and Engineering with a joint appointment at the Institute of Sustainable Transportation and Logics at the State University of New York at Buffalo. His research focuses on developing machine learning and signal processing tools to study the dynamics of large social systems in situ. He won the first Best Paper Award at the annual SBP conference with a work on predicting who gets you sniffle from your social network and the volunteer symptoms report. He has a PhD degree from the M.I.T. Media Laboratory. Email address is [wendong@buffalo.edu](mailto:wendong@buffalo.edu)