Neural Networks for Reinforcement Learning

Fall 2017

Instructor:	Emre Neftci	**Time:**	Fri 1:00PM – 3:50PM
Email:	eneftci@uci.edu	**Place:**	SBSG 2200.

Course Overview: Did you ever dream about a class where you could earn credits for playing games? You almost found it! This class will be about having fun by creating programs that use neural network to learn how to play games.

Neural networks in machine learning and reinforcement learning seem to be unstoppable in solving problems at proficiencies that were recently deemed reserved to humans (*e.g.* AlphaGo). What is the current status of deep reinforcement learning? Is a near future pervasive with artificial decision making systems a reality? What components are currently for missing reinforcement learning agents to become *bona fide* cognitive agents? This course will explore these questions through the current state-of-the-art in neural networks and reinforcement learning. The course will introduce and review "deep reinforcement learning" and use state-of-the-art software (PyTorch, OpenAI Gym) to replicate them in small virtual environments.

The instructor will start by introducing the reinforcement learning problem and the software tools necessary to run them (Python, PyTorch, OpenAI gym). One class will be devoted to teaching PyTorch and Neural Networks for classification and generation, including convolutional neural networks and generative adversarial networks. Note that this course will **not** study the theoretical aspects of the algorithms, but rather focus on their practice.

Using the Pytorch and OpenAI Gym framework, we will first solve reinforcement learning "classical control" problems (*e.g.* cartpole) using simple policy improvement algorithms. Next we will study deep Q-learning Networks (DQN), Actor Critic methods and Hierarchical reinforcement learning. We will use these methods to learn to play ATARI 7600 games. A guest lecture will introduce AlphaGo, and the applications of these algorithms for solving other problems. Finally, the instructor will provide an outlook on his own research in deep reinforcement learning, intrinsic motivation, and their instantiations in brain-inspired computers.

The first half of the class will be lecture based. The second half will involve the students present a paper chosen among a list prepared by the instructor. Every student will be required to apply the proposed method to a problem of their choice. The written software and a report will be used as a basis for the course assessment.

This course is ideal for students interested in reinforcement learning, and curious about the current state-of-the-art solutions in computers. The course is introductory and intended to be accessible to students from a broad range of disciplines, with varying background knowledge in the field. However, quantitative reasoning skills, including basic calculus and computer programming (Python) are necessary. In case of doubt, interested students are encouraged to e-mail the instructor.

The programs distributed for the hands-on exercises and assignments will be written in Python, using PyTorch and OpenAI gym packages. The instructor will set up a Jupyter Notebook server with all software pre-installed (A docker container will be made available for students willing to run on their own computers). The software engineering aspect will be kept at a minimum, and programming will involve deploying existing code on to new problems. If necessary, the instructor will introduce the basic Python programming during the second week. As a Python tutorial, refer to these resources: http://www.scipy-lectures.org/intro/index.html (sections 1.1 through 1.4).

The course will follow [1] and make use of their software. Additionally, the course will refer to the following books.

Key References:

[1] Maxim Lapan. *Deep Reinforcement Learning Hands-On: Apply modern RL methods, with deep Q-networks, value iteration, policy gradients, TRPO, AlphaGo Zero and more*. Packt Publishing Ltd, 2018.

[2] Richard S Sutton and Andrew G Barto. *Reinforcement Learning: An introduction (bibinfoedition2 ed.)* 2018.

[3] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. "Asynchronous methods for deep reinforcement learning". **International conference on machine learning**. 2016, pp. 1928–1937.

[4] David Silver et al. "Mastering the game of Go with deep neural networks and tree search". **Nature** 529.7587 (2016), pp. 484–489.

[5] Volodymyr Mnih et al. "Human-level control through deep reinforcement learning". **Nature** 518.7540 (2015), pp. 529–533.

Course Pages: `https://canvas.eee.uci.edu/courses/13699`

Teaching Assistant: None

Office Hours:

- **Neftci:** Wednesdays 1:00PM - 5:00 by appointment at SBSG 2308. Appointements can be booked through Google Calendar

Tentative Course Schedule:

The Reinforcement Learning Problem	Jan 11
Deep Learning with Pytorch	Jan 18
OpenAI Gym and Policy Network Agents	Jan 25
The Bellman Equation: Value Iteration and Policy Gradient Methods	Jan 25
Deep Q Networks	Feb 1
Actor Critic and A2C methods	Feb 8
Paper Presentations	Feb 15
Paper Presentations	Feb 22
Paper Presentations	Mar 1
AlphaGO and the APV-MCTS algorithm (Guest Lecture)	Mar 8
Paper Presentations Mar 15	

Grading: Assignments: (50%), Report (50%). Reports and assignments must be submitted before the deadline posted with each assignment sheet. There will no more than 4 individually graded assignments. The overall grade for the assignments will be calculated using the best 3 grades.

Important Dates:

Project Report Due	Friday after end of term