Course Syllabus

VE693

Deep Reinforcement Learning

Fall 2021

Course Description:

Reinforcement learning is a generic artificial intelligence approach for autonomously learning from trials and error. Deep reinforcement learning is the integration of deep learning and reinforcement learning. This course will cover the main recent algorithms and approaches in deep reinforcement learning and illustrate them on various applications (traffic light control, video games, robotics…).

This course provides a presentation of deep reinforcement learning (RL), with a focus on its recent developments. Topics include RL basics, deep Q-learning, policy gradients, actor-critic algorithms, model-based RL, imitation learning, inverse RL, hierarchical RL, and multi-agent RL.

Instructor:

Name: Paul Weng
Email: paul.weng@sjtu.edu.cn
Phone: +86-21-3420-6765 (Ext. 4061)
Office: Room 406
Office hour: Mon. 9am-11am

COURSE OBJECTIVES:

1. To provide an overview of deep RL as a field of study
2. To explain the fundamental issues and principles in deep RL
3. To cover the main deep RL approaches
4. To practice academic communication (reading and orally presenting research papers)

COURSE OUTCOMES:

1. To be able to recognize problems that can be solved by (deep) RL
2. To know how to apply RL techniques to solve practical problems, possibly after reformulating them such that they are amenable to such techniques

3. To become familiar with the main deep RL algorithms and know when and how to apply them
4. To be able to read research papers
5. To improve oral communication skills

Textbook (Author, Book Title, Publisher, Publication Year, ISBN)**:**

The course has no required textbook, but the following books could be used as a reference:

S. Russell and P. Norvig. Artificial Intelligence: a Modern Approach. Pearson.

I. Goodfellow, Yoshua Bengio, Aaron Courville. Deep Learning. MIT Press.

R.S. Sutton and A.G. Barto. Reinforcement Learning: an Introduction. MIT Press.

V. François-Lavet, P. Henderson, R. Islam, R., M.G. Bellemare, and J. Pineau. An Introduction to Deep Reinforcement Learning. Foundations and Trends in Machine Learning.

Course Prerequisites:

Necessary computer science and mathematical background covered in VE281, VE203, VV216/256/286

Course Website:

Lecture slides, assignments, and grades will be posted on the class webpage on Canvas:

https://umjicanvas.com/courses/2331

If you cannot attend the class physically, the online lectures and recordings can be accessed via the Feishu group of this course. The QR code to join this join is shared on Canvas. We will also use Piazza for announcements and discussions.

Grading Policy (Assignments %, Project, Exams, etc.)**:**

The students will be evaluated with:

1. Participation and etiquette (5%)
2. Online graded quizzes (10%)
3. Assignments (20%)
4. Final project (25%)
5. Mid-term exam (written) (20%)
6. Final exam (written) (20%)

Any questions about the grading of assignments or exams must be brought to the attention of the instructor within one week after the grade is released.

Honor Code Policy:

All students are expected to abide by the JI's Honor Code Policy. Although oral discussions with classmates

are encouraged, all assignments and exams are individual efforts. Therefore, reading or copying another person's solution is strictly forbidden. In all cases in which we have reason to believe that cheating has occurred, we will submit relevant materials to the Honor Council for evaluation.

Any course material (e.g., slides, projects, notebooks) should not be shared (especially on public websites) even after the course has ended. For the solutions you have written, please do not share them as well so that future students are not tempted to copy them.

Teaching Schedule:

Week	NO.	lectures and Exams	Comments
1	1	Introduction to reinforcement learning	
	2	Markov decision processes and variants	
2	3	Planning	
	4	Overview of machine learning	
3	5	Behavior cloning	
	6	Model-free prediction	
4	7	Model-free control	
	8	Value function approximation	
5	9	Policy gradient	
	10	Actor-critic methods	
6	11	Midterm review	
	12	Advanced value-based deep RL methods	
7	13	Midterm exam	
	14	Advanced policy-based deep RL methods	
8	15	Integrating model learning and policy learning	
	16	Inverse reinforcement learning	
9	17	Batch reinforcement learning, off-policy learning	
	18	Safe reinforcement learning	
10	19	Generalization, transfer learning	
	20	Multi-task learning	
11	21	Hierarchical reinforcement learning	
	22	Multi-agent reinforcement learning	
12	23	Final review	
	24	Final project presentation	
13	25	Final exam	
	26		