

EEC289A: An Introduction of Reinforcement Learning

Instructor: Prof. Lifeng Lai
Office: 3163 Kemper Hall
Phone: 530-752-7978
Email: lflai@ucdavis.edu

Course description: This course focuses on the introduction of one important subject of machine learning: reinforcement learning, which is considered the core for artificial intelligence. Topics include fundamentals of reinforcement learning, bandit problems, Markov decision processes, dynamic programming, Monte Carlo methods, temporal-difference learning, on-policy vs. off-policy learning, learning vs. planning, approximation methods, eligibility trace, policy gradient methods, and critic-actor methods.

Number of Units: 3
Hours per week: 3
Class number: EEC 289A

Topics:

- Fundamentals of reinforcement learning
- Bandit problems
- Markov decision process
- Dynamic programming
- Monte Carlo methods
- Temporal-difference learning: TD(0)
- N-step TD
- Learning vs. planning
- Approximation methods
- Eligibility trace
- Policy gradient methods

Recommended Textbook:

Reinforcement Learning: An Introduction, Second edition

By Richard S. Sutton and Andrew G. Barto
ISBN-10: 0262039249
ISBN-13: 978-0262039246

Supplementary Textbook:

Reinforcement Learning and Optimal Control

By Dimitri Bertsekas
ISBN-10: 1886529396
ISBN-13: 978-1886529397

Grading Structure:

1. Homework sets. 50%.

Homework sets will mainly be computer programming to implement learning algorithms discussed in the class.

2. Final project. 50%

The final project will focus on exploring the applications of learning algorithms discussed in the class.
