

EE-618

Theory and Methods for Reinforcement Learning

Cevher Volkan

Cursus	Sem.	Type
Electrical Engineering		Opt.

Language of teaching	English
Credits	3
Session	
Exam	Project report
Workload	90h
Hours	42
Courses	28
TP	14
Number of positions	30

Frequency

Every 2 years

Remark

Next time: Spring 2022

Summary

This course describes theory and methods for decision making under uncertainty under partial feedback.

Content

1. Introduction to the reinforcement learning (RL) paradigm
2. Overview of classical developments I: Markov Decision Process (MDP, POMDP), and Dynamic Programming (Value Iteration, Policy Iteration)
3. Overview of classical developments II: Monte-Carlo methods, TD-Learning, Q-Learning, SARSA (Model-based RL, and Model-free RL)
4. Stochastic Bandits and Thompson (posterior) Sampling
5. Bandit based RL algorithms (UCRL, UCAg, UCCRL, REGAL) - Exploration and Exploitation
6. Policy Search (Policy gradient algorithms, variance reduction, TRPO algorithm)
7. Imitation Learning (Inverse Reinforcement Learning, Apprenticeship Learning)

Keywords

Reinforcement learning, policy search.

Learning Prerequisites**Required courses**

Optimization, probability theory, mathematics of data.

Assessment methods

Project report.