EPFL

CS-456 **Artificial neural networks/reinforcement learning**

Gerstner Wulfram

Cursus	Sem.	Type
Biocomputing minor	E	Opt.
Computational Neurosciences minor	E	Opt.
Computational biology minor	E	Opt.
Computational science and Engineering	MA2, MA4	Opt.
Computer science	MA2, MA4	Opt.
Cybersecurity	MA2, MA4	Opt.
Data Science	MA2, MA4	Opt.
Data science minor	E	Opt.
Digital Humanities	MA2, MA4	Opt.
Electrical Engineering		Opt.
Financial engineering	MA2, MA4	Opt.
Life Sciences Engineering	MA2, MA4	Opt.
Neuro-X minor	E	Opt.
Neuro-X	MA2, MA4	Opt.
Quantum Science and Engineering	MA2, MA4	Opt.
SC master EPFL	MA2, MA4	Opt.
Statistics	MA2, MA4	Opt.

Language of teaching	English
Credits	6
Session	Summer
Semester	Spring
Exam	Written
Workload	180h
Weeks	14
Hours	**4 weekly**
Lecture	2 weekly
Exercises	2 weekly
Number of positions	

Summary

Since 2010 approaches in deep learning have revolutionized fields as diverse as computer vision, machine learning, or artificial intelligence. This course gives a systematic introduction into influential models of deep artificial neural networks, with a focus on Reinforcement Learning.

Content

- *General Introduction and Reinforcement Learning (RL) for Bandit Problems*

- *RL 1: Bellman equation and SARSA*
- *RL 2: Q-learning, n-step-TD learning, and eligibility traces*
- *RL 3: Continuous state space and function approximation*
- *RL 4: Policy gradient methods*
- *RL 5: Advantage Actor-Critic, eligibility traces, model-free/model-based*

- *Deep RL 1: Applications of Model-free RL in Video games and simulated Robotics*
- *Deep RL2: Applications iof Model-based RL: Board games and Replay buffer*
- *Deep rRL3: Markov Decision Processes and Policy iteration*

- *RL and the Brain: Three-factor Learning Rules*
- *RL and Hardware: Distributed algorithms and energy consumption,*
- *RL and Internal Rewards: Novelty and Surprise*
- *RL and Intrinsically Motivated Agents: Curiosity-driven Exploration*

Keywords

Deep learning, artificial neural networks, reinforcement learning, TD learning, SARSA, Actor-Critic Networks

Learning Prerequisites

Required courses

CS 433 Machine Learning (or equivalent)
Calculus, Linear Algebra (at the level equivalent to first 2 years of EPFL in STI or IC, such as Computer Science, Physics or Electrical Engineering)

Recommended courses

stochastic processes
optimization

Important concepts to start the course

- *Regularization in machine learning,*
- *Training base versus Test base, cross validation.*
- *Gradient descent. Stochastic gradient descent.*
- *Expectation, Poisson Process, Bernoulli Process.*

Learning Outcomes

By the end of the course, the student must be able to:
- Apply learning in deep networks to real data
- Assess / Evaluate performance of learning algorithms
- Elaborate relations between different mathematical algorithms of learning
- Judge limitations of algorithms
- Propose algoriothms and models for learning from experience
- Apply Reinforcement Learning

Transversal skills

- Continue to work through difficulties or initial failure to find optimal solutions.
- Access and evaluate appropriate sources of information.
- Write a scientific or technical report.
- Manage priorities.

Teaching methods

Ex cathedra lectures and miniproject.
Ex catherdra: Main ideas presented with slides and calculations presented on the blackboard. Every week the ex cathedra lectures are interrupted for one in-class exercise. The results of this exercise are needed for the second part of the lecture. Additional exercises are given as homework or can be disussed in the second exercise hour. Lectures are also interrupted by several short Quizzes.
Miniproject: The Miniprojects are done in a team of two and selected from a list of two or three miniprojects.

Expected student activities

Work on miniproject
Solve all exercises
Attend all lectures and take notes during lecture, participate in quizzes.
If you cannot attend a lecture, then you must read the recommended book chapters

Assessment methods

Written exam (70 percent) and miniproject (30 percent)

Supervision

Office hours	No
Assistants	Yes
Forum	Yes
Others	TAs are available during exercise sessions.
	Every week one of the exercises is run as 'integrated exercise' during the lecture.

Resources

Bibliography

•

• Textbook: Reinforcement Learning by Sutton and Barto (MIT Press). Pdfs of the preprint version of the book are availble online

Ressources en bibliothèque

• Reinforcement Learning / Sutton

Websites

• https://lcnwww.epfl.ch/gerstner/VideoLecturesANN-Gerstner.html
• https://moodle.epfl.ch/course/view.php?id=15633

Moodle Link

• https://go.epfl.ch/CS-456

Videos

• https://lcnwww.epfl.ch/gerstner/VideoLecturesANN-Gerstner.html