

EE-559

Deep learning

Cavallaro Andrea

Cursus	Sem.	Type
Civil & Environmental Engineering		Opt.
Computational science and Engineering	MA2, MA4	Opt.
Computational science and engineering minor	E	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.
Electrical and Electronical Engineering	MA2, MA4	Opt.
Financial engineering	MA2, MA4	Opt.
Learning Sciences		Opt.
Life Sciences Engineering	MA2, MA4	Opt.
Minor in Quantum Science and Engineering	E	Opt.
Neuro-X minor	E	Opt.
Neuro-X	MA2, MA4	Opt.
Quantum Science and Engineering	MA2, MA4	Opt.
Robotics, Control and Intelligent Systems		Opt.
Robotics	MA2, MA4	Opt.
SC master EPFL	MA2, MA4	Opt.
Statistics	MA2, MA4	Opt.

Language of teaching	English
Credits	4
Withdrawal Session	Unauthorized Summer
Semester Exam	Spring During the semester
Workload Weeks	120h 14
Hours	4 weekly
Courses	2 weekly
Exercises	2 weekly
Number of positions	150

It is not allowed to withdraw from this subject after the registration deadline.

Summary

This course explores how to design reliable discriminative and generative neural networks, the ethics of data acquisition and model deployment, as well as modern multi-modal models.

Content

This course equips students with a comprehensive foundation of modern deep learning, enabling students to design and train discriminative and generative neural networks for a wide range of tasks. Topics include:

- Deep learning applications (natural language processing, computer vision, audio processing, biology, robotics, science), principles and regulations
- Loss functions, data and labels, data provenance
- Training models: gradients and initialization
- Generalization and performance
- Transformers
- Graph neural networks
- Generative adversarial networks
- Variational autoencoders
- Diffusion models
- Multi-modal models
- Interpretability, explanations, bias and fairness

Keywords

machine learning, neural networks, deep learning, python

Learning Prerequisites

Required courses

- Basics in probabilities and statistics
- Linear algebra
- Differential calculus
- Python programming

Recommended courses

- Basics in optimization
- Basics in algorithmic
- Basics in signal processing

Important concepts to start the course

Discrete and continuous distributions, normal density, law of large numbers, conditional probabilities, Bayes, PCA, vector, matrix operations, Euclidean spaces, Jacobian, Hessian, chain rule, notion of minima, gradient descent, computational costs, Fourier transform, convolution.

Learning Outcomes

By the end of the course, the student must be able to:

- Interpret the performance of a deep learning model
- Analyze the limitations of a deep learning model
- Justify the choices for training and testing a deep learning model
- Propose new solutions for a given application

Transversal skills

- Respect relevant legal guidelines and ethical codes for the profession.
- Take account of the social and human dimensions of the engineering profession.
- Design and present a poster.
- Make an oral presentation.
- Demonstrate the capacity for critical thinking

Teaching methods

Ex-cathedra lectures, class discussion, exercises (using python), group project.

Expected student activities

Attendance to lectures, participation in discussions, completing exercises, completing a project, reading written material (scientific papers and books).

Assessment methods

Exercices and group project.

Resources

Références suggérées par la bibliothèque

- [Deep learning](#), Goodfellow, MIT Press, 2016

Notes/Handbook

Not mandatory: <http://www.deeplearningbook.org/>

Moodle Link

- <https://go.epfl.ch/EE-559>