EE-559 **Deep learning**

Fleuret François

Cursus	Sem.	Type
Computational science and Engineering	MA2, MA4	Opt.
Computer science	MA2, MA4	Opt.
Cybersecurity	MA2, MA4	Opt.
Data Science	MA2, MA4	Opt.
Digital Humanities	MA2, MA4	Opt.
Electrical Engineering		Obl.
Electrical and Electronical Engineering	MA2, MA4	Opt.
Life Sciences Engineering	MA2, MA4	Opt.
Robotics	MA2, MA4	Opt.
SC master EPFL	MA2, MA4	Opt.

Language of teaching	English
Credits	4
Session	Summer
Semester	Spring
Exam	Written
Workload	120h
Weeks	14
Hours	**4 weekly**
Courses	2 weekly
Exercises	2 weekly
Number of positions	**342**

Summary

The objective of this course is to provide a complete introduction to deep machine learning. How to design a neural network, how to train it, and what are the modern techniques that specifically handle very large networks.

Content

The course aims at teaching the required skills to use deep learning methods on applied problems.
It will show how to design and train a deep neural network for a given task, and the sufficient theoretical basis to go beyond the topics directly seen in the course.
The planned content of the course:

- What is deep learning, introduction to tensors.

- Basic machine-learning, empirical risk minimization, simple embeddings.

- Linear separability, multi-layer perceptrons, back-prop.

- Generalized networks, autograd, batch processing, convolutional networks.

- Initialization, optimization, and regularization. Drop-out, activation normalization, skip connections.

- Deep models for Computer Vision.

- Analysis of deep models.

- Auto-encoders, embeddings, and generative models.

- Recurrent models and Natural Language Processing.
- pytorch tensors, deep learning modules, and internals.

Concepts will be illustrated with examples in the pytorch framework (http://pytorch.org).

Keywords

machine learning, neural networks, deep learning, computer vision, python, pytorch

Learning Prerequisites

Required courses

- Linear algebra (vector, matrix operations, Euclidean spaces).
- Differential calculus (Jacobian, Hessian, chain rule).
- Python programming.
- Basics in probabilities and statistics (discrete and continuous distributions, normal density, law of large numbers, conditional probabilities, Bayes, PCA)

Recommended courses

- Basics in optimization (notion of minima, gradient descent).
- Basics in algorithmic (computational costs).
- Basics in signal processing (Fourier transform, wavelets).

Teaching methods

Ex-cathedra with exercise sessions and mini-projects. Invited speakers from the industry will present how deep learning is used in practice for their applications.

Assessment methods

Two mini-projects by groups of three students, and one final written exam.

Resources

Notes/Handbook

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