UNIVERSITÀ
DEGLI STUDI DI MILANO-BICOCCA

COURSE SYLLABUS

Foundations of Deep Learning

2425-1-FDS01Q012

Aims

The aim of this course is to provide the theoretical foundations of mathematics and statistics for deep learning including linear algebra, optimization, regularization, and dimensionality reduction. The most important deep neural network architectures will be covered in this course. Thanks to a practical part of the course, the student will be able to handle the main tools for deep learning and then design and optimize a deep neural network.

Contents

The course consists of a theoretical part and a practical part of laboratory. The theoretical part aims at exploring applied math, machine learning basics and deep neural networks. The practical part consists in basic and advanced exercises using deep learning frameworks.

Detailed program

- Introduction to Deep Learning: Overview of the field, its impact, and key concepts.
- Linear Algebra: Basic mathematical tools for deep learning including vectors, matrices, and linear algebra.
- Machine Learning Basics: Foundations of machine learning, types of learning, and introduction to the main algorithms.
- Gradient-Based Learning and Backpropagation: Understanding the mechanism of training neural networks through gradients and backpropagation.
- Feed Forward Networks: Multilayer Perceptron Networks, Activation functions.
- Convolutional Neural Networks (CNNs): Exploring relevant CNNs and their applications in processing visual data.
- Regularization for Deep Learning: Techniques to improve network generalization and prevent overfitting.

- Recurrent Neural Networks (RNNs): Detailed study of networks designed to process sequential data such as RNNs and LSTMs.
- Deep Transformers: Exploration of transformer models, focusing on architectures like GPT and BERT that transform the NLP landscape.
- Dimensionality Reduction Techniques: Methods like PCA and t-SNE for reducing the complexity of data while preserving its essential patterns.
- Autoencoders: Study of autoencoders for unsupervised learning and feature extraction.
- Practical Methodology: Best practices in applying deep learning including data preprocessing, model selection, training strategies, and evaluation.

Learning Outcomes:
By the end of this course, students will be able to:

- Understand and apply foundational mathematics for deep learning.
- Design and implement various types of neural networks.
- Evaluate and improve deep learning models using advanced techniques.
- Apply deep learning to practical problems across different domains.

Prerequisites

A basic understanding of calculus, statistics, and programming is recommended for students enrolling in this course.

Teaching form

The teaching includes a part of theoretical lectures that will be held in the classroom, and a part of laboratory that will be held in the laboratory and/or classroom and will require the use of one's own PC. The two parts will be based both on delivery mode and interactive mode.

Textbook and teaching resource

Main textbook:

- Goodfellow, I., Bengio, Y., Courville, A., & Bengio, Y. (2016). Deep learning (Vol. 1, No. 2). Cambridge: MIT press
  https://www.deeplearningbook.org/ https://github.com/janishar/mit-deep-learning-book-pdf
- Simon J.D. Prince (2023) Understanding Deep Learning Published by MIT Press Dec 5th 2023. https://udlbook.github.io/udlbook/

Additional Resources

- Neural Networks and Deep Learning by By Michael Nielsen (2016) – online book http://neuralnetworksanddeeplearning.com/

From the teacher

- Scientific articles suggested by the teacher (readings) Teachers' slides (http://elearning.unimib.it/)

Semester

Second Semester

Assessment method

The exam consists in two parts:

- The first one is an individual test. It is about code comprehension on basic machine learning problems and some theoretical questions on the first part of the course (this first part of the exam can be taken in April, and in every exam starting from June).
- The second part of the exam is the design and realization of a project assigned by the teacher about deep learning classification or recognition or regression. The project can be developed individually or in collaboration with one colleague. The project will be discussed as oral presentation and the teachers can ask questions about the second theoretical and practical parts of the course (this part of the exam can be taken starting from June).

Office hours

Upon appointment,
Paolo Napoletano, Monday from 14 to 16
Marco Buzzelli, Monday from 14 to 16

Sustainable Development Goals