UNIVERSITÀ
DEGLI STUDI DI MILANO-BICOCCA

COURSE SYLLABUS

Advanced Artificial Intelligence, Machine Learning and Deep Learning

2627-1-F9103Q009

Aims

Recent developments in Artificial Intelligence and Machine Learning have changed how we work and live by assuming a significant role in the industry, education, and, more generally, in today's culture. This course provides students with advanced techniques for designing intelligent systems, understanding model behavior, learning effective representations, transferring knowledge across tasks, and reasoning with modern neural architectures. The course combines theoretical foundations with practical applications.

Contents

The course is organized around four main themes:

1. **Explainable Artificial Intelligence (XAI)**: understanding and explaining the decisions of modern machine learning models.
2. **Representation Learning**: learning useful data representations through embeddings, neural models, and graph-based approaches.
3. **Attention, and Knowledge Representation**: processing information through attention mechanisms and reasoning over structured knowledge.
4. **Transfer Learning**: reusing learned knowledge across tasks and domains to improve learning efficiency and performance.

Detailed program

A) Explainable Artificial Intelligence (XAI)

- Interpretability and explainability
- Foundations of XAI
- Post-hoc explanation methods
- LIME (Local Interpretable Model-Agnostic Explanations)
- Saliency Maps for tabular data
- Local and global explanations

B) Representation Learning

- Neural embeddings and latent spaces
- Representation learning principles
- Skip-Gram and word embeddings
- Graph representation learning
- Introduction to Knowledge Graphs
- Knowledge graph embeddings (TransE)

C) Attention and Knowledge-Based Reasoning

- Knowledge representation and reasoning
- Knowledge Graphs and relational learning
- Translational embedding models
- Attention mechanisms
- Query-Key-Value framework
- Self-attention
- Attention in modern AI systems and Transformers

D) Transfer Learning

- Foundations of transfer learning
- Inductive transfer learning
- Transductive transfer learning
- Unsupervised transfer learning
- Domain adaptation and domain alignment
- Representation transfer
- Practical applications of transfer learning

Prerequisites

Most of the prerequisites will be briefly recalled in classes. However, basic knowledge of Linear algebra, Calculus, and Probability are warmly recommended. Basic programming skills are fundamentals. Python programming language is strongly recommended.

Teaching form

Lectures introduce the main concepts, models, and algorithms. Practical sessions and laboratories allow students to apply the techniques discussed in class through guided exercises and small projects. Active participation in both lectures and labs is strongly encouraged.

Textbook and teaching resource

A) Suggested texts (Specialized papers and further resources will be provided during the course)

- Zhang, Aston, et al. "Dive into deep learning." arXiv preprint arXiv:2106.11342 (2021).
- Murphy, Kevin P. Machine learning: a probabilistic perspective. MIT press, 2012.
- Heaton, Jeff. "Ian Goodfellow, Yoshua Bengio, and Aaron Courville: Deep learning." (2017).
- Russell, Stuart J. Artificial intelligence a modern approach. Pearson Education, Inc., 2010. (and subsequent editions)

Semester

Second

Assessment method

The final grade will be based on:

- Project work and presentation
- Written examination consisting of theoretical questions and simple exercises covering the course topics

Office hours

By appointment.

Sustainable Development Goals

QUALITY EDUCATION