



UNIVERSITÀ
DEGLI STUDI DI MILANO-BICOCCA

SYLLABUS DEL CORSO

Artificial Intelligence

2223-2-F1801Q155

Aims

The aim of the course is to enable the student to master basic knowledge and tools necessary to understand, use and develop systems based on Artificial Intelligence (AI), together with the ability to analyze classes of problems particularly suitable for being treated with methods and techniques characterizing the discipline.

Artificial Intelligence has developed by addressing a wide range of problems, from natural language processing to coordination of autonomous agents, thus branching out into a broad set of methodological approaches and more specific disciplines such as knowledge representation, multi-agent systems, machine learning, robotics, etc .. The more specific objective of this course is therefore twofold:

- provide an **integrated framework for the discipline based on the concept of intelligent agent**, with the aim of enabling the student to frame different methodological approaches and specific disciplines in the context of the development of intelligent systems.
- provide an in-depth study of **methodologies based on knowledge representation and learning**, with particular attention to the integration of paradigms based on symbolic AI (knowledge representation and reasoning) and based on machine learning (especially based on deep neural networks), with the aim of providing students with skills useful for developing knowledge-based systems that integrate the two paradigms.

On the one hand, we intend to provide an overall picture about AI today, also indicating relationships with the contents provided in other courses; on the other hand, we intend to present specific models and techniques that are already essential building blocks of important applications of Artificial Intelligence in science and industry (knowledge graphs, inference engines, embeddings, natural language processing).

Finally, considering the differences but also increasingly close relationships between scientific research and industrial applications, the course intends to discuss the limits of existing methodologies, both from the scientific point of view and from the point of view of industrial applications, and the consequent challenges for the future of AI.

Contents

- Introduction to modern AI and agents as unifying abstractions for intelligent systems
- Knowledge representation and reasoning: how to encode knowledge so as to support robust downstream AI applications
- Principles of representation learning with deep neural networks: how to learn representations from data.
- Building knowledge bases from unstructured data
- Deep integration of representation, reasoning and learning: neuro-symbolic integration

Detailed program

1. Introduction: learning and reasoning in Artificial Intelligence; interpretation, reasoning, prevision and control; autonomous agents (definition, classification, behavior, agent-bases models with simple reflexes, with memory, goal-based and utility-based).
2. Knowledge representation and reasoning: knowledge graphs, reasoning with ontologies (RDFS, OWL); rule-based and non monotonic reasoning (datalog, Answer Set Programming).
3. Knowledge representation and deep learning: distributional semantics, representation in vector spaces; representation learning (knowledge graph embeddings, link prediction, alignment between representation spaces).
4. Introduction to information extraction from unstructured data: named entity recognition, entity linking, relation extraction.
5. Introduction to neuro-symbolic integration: Logic Tensor Network and KENN.

Prerequisites

Basic knowledge of logics and mathematics. Basic knowledge of machine learning.

Teaching form

Lectures and exercise with students' personal computers. Moodle e-learning platform. Seminars about applications by experts from the industry.

The course is taught in English.

Textbook and teaching resource

Textbooks (specific sections):

- Artificial Intelligence: a Modern Approach. S.J. Russell, P. Norvig, 4th Edition, Pearson, 2020.
- Artificial Intelligence: Foundations of Computational Agents, second edition. Poole, D, and Mackworth, A. Cambridge University Press 2017.
- Knowledge Graphs: Fundamentals, Techniques, and Applications. Kejriwal, Mayank, Craig A. Knoblock, and Pedro Szekely. MIT Press, 2021.

More material provided during the course.

Semester

Semester I

Assessment method

The final evaluation consists of the aggregation of the scores obtained in two independent assessments:

- The first assessment is based on an **exam-tailored project**, carried out individually or in groups and aimed at bringing the student to have in-depth knowledge and/or hands-on experience of a specific topic covered in the course or linked to topics covered in the course; the project is discussed through an oral presentation supported by slides lasting about 20 minutes; it is possible, during the presentation, to include a short demo of the project. The evaluation is based on: significance of the project for the topics covered in the course, methodological soundness (within the limits of what is reasonable to ask for an exam project); mastery of the in-depth topic demonstrated during the oral presentation.
- The second assessment is based on the **evaluation of the knowledge acquired by the student on the topics addressed during the course** through the discussion of assignments that students must execute individually as homework. Assignments will be evaluated and discussed during the oral exam after the presentation of the project.

Office hours

On demand.

Sustainable Development Goals

QUALITY EDUCATION

