UNIVERSITÀ
DEGLI STUDI DI MILANO-BICOCCA

COURSE SYLLABUS

Artificial Intelligence

2425-2-F1801Q155

Aims

The objective of the course is to enable students to master the basic knowledge and tools necessary to understand, use, and create Artificial Intelligence (AI) systems, along with the ability to analyze different classes of problems and solutions based on AI. A special focus throughout the course will be dedicated to the intersection of machine learning, with special consideration of Large Language Models (LLMs), and multi-agent and knowledge-based systems, taking a perspective that is frequently referred to as neuro-symbolic AI. ,

The course has a dual objective:

- Provide an overview of the discipline focused on the concept of intelligent agent, enabling students to critically frame problems, solutions, and specific methodological approaches within the development of intelligent systems.
- Provide an in-depth exploration of some themes and solutions of particular importance for the development of intelligent systems today, namely: 1) autonomous AI, 2) embodied and affective AI, 3) knowledge-based AI.

Contents

The first part of the course (around 1/2 of the course) will cover four foundational topics:

- The **different paradigms** proposed in AI and the challenge of ethics in AI
- The **agent-based paradigm** as a conceptual model to organically frame various problems and models proposed in modern Artificial Intelligence. Agents will be characterized from the perspective of autonomy and relationships such as those between agent and environment, behavior and perception, and behavior and knowledge.
- **Knowledge graphs** as abstractions to support a variety of knowledge-based applications where reasoning,

question answering and data linking play a crucial role. Since topics such as RDF, Ontologies and reasoning frameworks like RDFS and OWL are now covered in BA courses, in the course we will quickly recap these concepts.
- **Distributional semantics and large language models** as approaches to learning representations from data, understanding input data, and generating novel content

The second part of the course (around 1/2 of the course) will focus on more specific topics:

- **Conversational agents and related interaction models**. This part is dedicated to modeling the behavior of intelligent agents based on coordination in multi-agent systems. Models for simulation through multi-agent systems, perception, and interaction modeling in complex systems, will be discussed. The applications covered in this part pertain to simulation systems and systems based on the adaptive behavior of agents.
- **NLP and LLMs for Knowledge Graph-powered AI systems.** This part covers two sub-topics: information extraction techniques and LLMs to support data enrichment and knowledge base construction; approaches combining generative models and knowledge bases such as Retrieval Augmented Generation.
- **Representation learning for Knowledge Graphs and neuro-symbolic integration**, exploring machine learning methods to build soft inference systems on top of Knowledge Graphs and other structured knowledge representation frameworks, covering Knowledge Graph embeddings and deep neurosymbolic integration frameworks like Logic Tensor Networks.

Detailed program

1. Introduction: learning and reasoning in Artificial Intelligence; interpretation, reasoning, prediction, control; the concept of autonomous agent (definition, classification, behavior, models of agents with simple reflexes, with memory, goal-based, utility-based).
2. Models and mechanisms of interaction in Multi-Agent Systems (MAS): collective artificial intelligence and complex systems; modeling, simulation, and analysis of self-organizing behaviors.
3. Conversational agents: modeling, simulation, analysis of self-organizing behaviors; integrating LLMs in Multi-agent Systems.
4. From knowledge representation to representation learning: knowledge bases, knowledge graphs, and ontologies; learning linguistic representations, word embeddings, and large language models (BERT, GPT-X); reasoning in LLMs; risks and ethical issues in LLMs
5. Knowledge base construction and exploitation with LLMs: information extraction from texts (Named Entity Recognition, Named Entity Linking, introduction to relation extraction); information extraction from tabular data (tabular data understanding and annotation); Retrieval Augmented Generation with knowledge bases and annotated data.
6. Learning representations of knowledge graphs using neural networks: knowledge graph embeddings; AI frameworks based on neuro-symbolic integration (Logic Tensor Network).

Prerequisites

Basic knowledge of logics and mathematics. Basic knowledge about machine learning and deep neural networks.

Teaching form

Lectures and exercise with students' personal computers. Moodle e-learning platform. Seminars about usage of semantics in real-world applications given by experts from the industry.

The course is teached in English.

Textbook and teaching resource

Textbooks :

S.J. Russell, P. Norvig, “Intelligenza Artificiale: un approccio moderno“, 2a edizione, Pearson - Prentice Hall, 2005 (volume 1)

J. Ferber, Multi-agent systems: An introduction to distributed artificial intelligence, Addison-Wesley Professional, 1999: sintesi a dispense disponibile sull’e.learning del Corso.

Recommended reading:

C. Cornoldi, L'intelligenza, Il Mulino Ed., 2009. Cesare Cornoldi. Formicai, imperi, cervelli: introduzione alla scienza della Complessità (Il Mulino, 2007), pp. 235.

Grigoris Antoniou, Paul Groth, Frank van Harmelen, Rinke Hoekstra. A Semantic Web Primer (Information Systems) third edition. The MIT Press; third edition edition (August 24, 2012), pp. 288.

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 or a survey, carried out individually or in groups, and aimed at bringing the student to have an 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 and the survey are both 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 survey consists of a bibliographic review on a topic, in which the student discusses and compares proposed solutions in the state of the art to a specific problem of interest for him. The evaluation is based on: significance of the project with respect to 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 verification of the knowledge acquired by the student about the topics addressed during the course by the execution of assignments related to these topics.

Office hours

On demand.

Sustainable Development Goals

QUALITY EDUCATION