SYLLABUS DEL CORSO

Artificial Intelligence

2021-2-F1801Q155

Aims

The aim of the course is to enable the student to master the knowledge and basic tools necessary to comprehend, use and create Artificial Intelligence systems, together with the ability to analyze classes of problems particularly suitable to be solved with methods and techniques that characterize the discipline. Because of the practical and at the same time explorative nature of Artificial Intelligence, the course has the objective of discussing how to discriminate between solutions that can be applied in industrial settings and innovative directions of research in this field. The students will be given conceptual, computational and basic methodological tools to understand and develop innovative solutions to automation problems through advanced Artificial Intelligence techniques. The course is aimed at students who aspire to enter into work and research environments where innovative choices prevail for the solution of complex problems and areas with a strong multidisciplinary component.

Contents

The course leverage the Agent-based Paradigm as a conceptual tool to frame and connect different problems and computational models addressed by contemporary Artificial Intelligence. Using this paradigm as a reference, several prominent aspects of knowledge processing will be vertically discussed duirng the course: interaction, coordination and collective intelligence; knowledge representation and reasoning; knowledge acquisition and perception; communication and harmonization between different representations; the deep relationthip between learning and reasoning.
The topics addressed in the course are based on a selection of computational models that are deemed impactful on a large number of innovative technological scenarios of today and tomorrow that requires complex knowledge-based solutions.

The first, shorter, part of the course will deal with the agent-based AI Paradigm, Multi-agent systems (MAS) and collective AI; in this part of the course, students will be introduced to models whose intelligent overall behavior is

resulting from the action and interaction of agents that are not necessarily characterized by a sophisticated individual behavior. Exemplary applications of this paradigm will be discussed in relation to the simulation of complex systems, self-organization and ambient intelligence. During the classes the presentation and critical discussion of advanced AI based real projects and paradigmatic case studies will be provided, in order to introduce also the multidisciplinary scenario of the faced problems. The second part of the course will deal with intelligence as the ability of an agent to elaborate symbolic information, to link it together and to infer new knowledge using automatic reasoning methods; in this part of the course we will study semantic models for the acquisition, representation and processing of knowledge at a large scale, with reference to innovative techniques now part of the most advanced ICT technologies on the market (think of systems such as IBM Watson and the Google Knowledge Graph). Particular emphasis will be devoted to automatic learning techniques applied to the representation of knowledge and its exploration. Finally, we will discuss one of the most hot direction of contemporary Artificial Intelligence, i.e., the native integration of machine learning and knowledge representation technologies, often referred to as 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, ontologies, inference rules (RDFS, OWL, SWRL).

3. Introduction to information extraction from unstructured data: named entity recognition, entity linking, relation extraction.

4. Knowledge representation and deep learning: distributional semantics, representation in vector spaces; representation learning (knowledge graph embeddings, link prediction, alignment between representation spaces); neuro-symbolic integration and open problems.

5. Models and interaction mechanisms in multi-agent systems (MAS); collective Artificial Intelligence and complex systems; modeling, simulation, and analysis of self-organizing behavior; reinforcement learning.

Prerequisites

Basic knowledge of logics and mathematics.

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.

During the Covid-19 emergency period, lessons will take place in a mixed mode: partial attendance and

asynchronous videotaped lessons. If this mode is not possible, the course will be held remotely with asynchronous videotaped lessons and synchronous events (QA sessions) in videoconference.

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.

Tommaso Di Noia, Roberto De Virgilio, Eugenio Di Sciascio, Francesco M. Donin. Semantic Web. Tra ontologie e Open Data. 1° ed. (Apogeo, 2013), pp. 240

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 II

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 in one of the following ways, freely chosen by the student:

1. oral exam taken in conjunction with the discussion of the first evaluation;
2. two ongoing tests consisting of exercises and open questions: one related to the topics covered in the first part of the course (autonomous agents, multi-agent systems and interaction), and one related to the topics covered in the second part of the course (knowledge representation, information extraction, models for knowledge exploration, representation learning).

In the Covid-19 emergency period, if it is not possible to carry out written exams in the presence, the ongoing tests will be replaced by assignments to be carried out individually and verified during the oral exam.

During the Covid-19 emergency period, oral exams will be online only.
They will be hosted using the WebEx platform and a public link will be published on the e-learning page to grant access to virtual spectators.

Office hours

On demand.