Computer Sciences, CS, CUNY

CSc 74011 Artificial Intelligence (3 hours; 3 credits)

Syllabus- Full 2020

Instructor: Sos Agaian, Ph.D. (Math) and Ph.D. (Engineering Sciences), Distinguished Professor

Office Hours: Fraidy-10:45am–11:45 pm or by appointment. If you do need to reach me, the best way is to come to my office hours. The next best way is by e-mail. On the other hand, please be aware that I receive a large volume of student e-mails, so I will not be able to respond right away)
Lecture: Friday 11:45am–1:45
Phone: 718-982-2850
E-mail: sos.agaian@csi.cuny.edu

Rationale

Artificial intelligence (AI) is a research field that studies how to realize intelligent human behaviors on a computer. The fundamental goal of AI is to make a computer that can learn, plan, and solve problems independently. This course aims to give an overview of some basic AI algorithms and an understanding of the possibilities and limitations of AI.

Course description

This is an introductory, graduate-level course on artificial intelligence. It emphasizes fast and smart search heuristics, thoughtful ways to represent knowledge, and incisive techniques that support rational decision-making. Application areas will include computer vision, natural language processing, and robotics. Other topics will be covered as time permits.

Prerequisites

Students are expected to have a solid background in the analysis of algorithms, proofs in propositional and first-order logic, discrete mathematics, and elementary probability.

Learning Objectives and Outcomes

The primary purpose of this course is to provide the most fundamental knowledge to the students so that they can ***understand what the AI*** is and ***how to use it in practice.***

Students who complete this course will be able to:

- Formulate search problems and implement search algorithms using admissible heuristics.
- Describe the state-space search as a mechanism for problem-solving, including optimal solutions and their complexity.
- Define machine learning and describe the specifics of several prominent machine-learning methods (e.g., SVMs, decision trees, Bayes nets, and artificial neural networks)
- Describe and illustrate the role of constraint satisfaction in AI, with appropriate examples.
- Discuss the role of probabilistic reasoning and mechanisms that employ it
- Apply selected basic AI techniques; judge applicability of more advanced techniques.
- Design, implement and evaluate a computer-based system, process, component, or program to meet desired needs
- Applications of AI (Natural Language Processing, Robotics/Vision)

- Get hands-on experience by solving real-world AI problems.

Course reading: Stuart Russell and Peter Norvig, Artificial Intelligence: A Modern Approach, Third Edition 2010. Pearson Education, Inc. ISBN: 978-0-13-604259-4 (http://aima.cs.berkeley.edu/). Students will also be required to read a wide variety of assigned papers and summarize and react to their content.

Also recommended: Andreas C. Müller, Sarah Guido, Introduction to Machine Learning with Python: A Guide for Data Scientists 1st Edition, ISBN-13: 978-1449369415 ISBN-10: 1449369413

Grades will be based on:

- Class participation 10%
- Assignments 60%
- Term project (presentation and report) 30%

Students must complete several projects. Projects will include hands-on application of basic AI techniques as well as the selection of appropriate technologies for a given problem. In a final project, groups of students will participate in the creation of an AI-based application to solve real-world issues including, search, computer vision, machine learning, logic, and constraint satisfaction problems. The project grade based on three aspects: project concept, results, presentation quality, and report quality. The final project will allow you to implement the skills you learned in the masters of AI.

List of topics

- **Week 1: Introduction to AI, history of AI, course structure and policies and present some AI applications**
  - What is AI?
  - AI now: Why this time it's different
  - Intelligent Agents
- **Week 2: Problem Solving**
  - Solving problems by searching
  - Knowledge and rationality
  - Heuristic search strategies
  - Adversarial search
  - Search and optimization (gradient descent)
- **Week 3: Knowledge Representation (**Acquire and represent knowledge about a domain**) and Reasoning(**Use the knowledge to solve problems in that domain**)**
  - Logical agents
  - First-order logic
  - Uncertain and probabilistic reasoning (Bayesian reasoning)
  - Bayesian networks: representation, independence, and inference
  - Hidden Markov model
- **Week 4: Learning: Supervised methods**
  - What is learning?
  - Learning from examples
  - Knowledge of Learning
  - Supervised vs. Unsupervised learning
  - Regression -- linear, logistic, ridge
  - Classification – decision trees, SVM

    - Model performance evaluation
    - Case study: Computer vision, text related supervised methods
- **Week 5: Learning-Unsupervised Methods**
    - Dimensionality reduction: PCA
    - Clustering – k-means, hierarchical clustering
    - Semi-supervised methods
    - Reinforcement learning
    - Case study: Computer vision, text clustering
- **Week 6: Deep Learning**
    - Neural networks and back-propagation
    - Convolutional neural networks
    - Recurrent neural networks
- **Week 7-8: AI applications (Computer Vision)**
    - Introduction to computer vision
    - Introduction to feature engineering
    - Object classification
    - Use of pre-trained models (Inception)
- **Week 9: AI applications (Natural Language Processing)**
- **Week 10: Current Directions in AI Research**
    - Case study: Uber and Facebook
- **Week 11: AI final project presentations**
- **Week 12: Review and conclusion**

Syllabus Change Policy: This syllabus may be subject to change with reasonable advanced notice.

Course Material

- You are responsible for all the material in the assigned readings, regardless of whether it has been explicitly covered in class.
- You are also responsible for all the material covered in class.
- It is strongly recommended that you read the assigned readings/ before each class. It will help you understand the material better when I lecture

Long-term goal: empirical research

Empirical AI research addresses a real-world problem with appropriate knowledge representations and a reasoning methodology for it, identifies or constructs algorithms to address it, and implements, tests, and evaluates alternative solution(s) to it. This course is intended to provide a solid foundation for empirical AI research**.**