New Jersey Institute of Technology
Department of Computer Science

CS670 - Artificial Intelligence - Spring'2023

Monday 2:30 - 5:20 PM, KUPF 106

Course Description | Outcomes | Readings | Tentative Contents | Grading Policy | Miscellaneous

---

Chengjun Liu, Ph.D.

Email: *cliu@njit.edu*
Phone: 973-596-5280
Office: GITC 4306
Office Hours: Monday 1:20-2:20PM & Friday 1:20-3:00PM or by appointment

---

Course Description

- This course introduces concepts, approaches and techniques of artificial intelligence, and focuses on materials that are fundamental and have a broad scope of applications. Topics include Problem Solving, Intelligent Agents, Logical Agents, Propositional Logic Knowledge Representation and Reasoning, First-Order Logic Knowledge Representation and Reasoning, Uncertain Knowledge Representation and Reasoning, Quantifying Uncertainty, Probabilistic Reasoning, Learning, Statistical Learning Theory, Bayesian Learning, Decision Tree learning, Neural Networks, Deep Learning, Reinforcement Learning, Genetic Algorithms, Perception, Pictorial Knowledge Representation, Search, Action, and Robotics.

- Prerequisite: CS 610 – Data Structures and Algorithms

Measurable Learning Outcomes

- Students learn the concepts, approaches and techniques of artificial intelligence.
- Students learn the materials that are fundamental and have a broad scope of applications in artificial intelligence, such as Problem Solving, Intelligent Agents, Knowledge and Reasoning, Logical Agents, Propositional Logic, First-Order Logic, Uncertain Knowledge and Reasoning, Quantifying Uncertainty, Probabilistic Reasoning, Learning, General Learning Model, Decision Tree learning, Unsupervised Learning, Supervised Learning, Statistical Learning Theory, Structural Risk Minimization, Support Vector Machine, Perception.

Readings

- S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, 4th edition, Prentice Hall, 2020.
- V. N. Vapnik, *The Nature of Statistical Learning Theory*, 2nd edition, Springer, 2000.
- Selected papers and handouts.

Tentative Contents

1. Introduction
    - AI Fundamentals (Turing test, cognitive science, logic, learning, games, robot, vehicle, agent)
    - AI prehistory and AI history: connectionism, symbolism, AI winters
    - Programming Languages: Lisp, Prolog, Matlab, C/C++, Java
    - Related Fields: ML, NN, EC, CV, PR, IP
2. Problem Solving
    - Intelligent Agents
    - Solving Problems by Searching
    - Breadth-first Search, Depth-first Search

    - Best-first Search, Greedy Search, A* Search
    - Games (Adversarial Search, Alpha-Beta Pruning)
3. Knowledge and Reasoning - Logical Agents
    - Knowledge-Based Agents
    - Logic, Propositional Logic
    - Models, Semantics, Inference, Validity and Satisfiability
    - Propositional Theorem Proving, Resolution, CNF
    - Games (Wumpus World)
4. Knowledge and Reasoning - First-Order Logic
    - FOL Syntax and Semantics
    - FOL Sentences, Models, Interpretation
    - FOL Quantification, Properties of Quantifiers
    - FOL KBs, Deducing Hidden Properties
5. Knowledge and Reasoning - Inference in First-Order Logic
    - Propositional vs. First-Order Inference
    - Universal and Existential Instantiation
    - Unification, GMP, Soundness of GMP
    - FOL KB and Resolution
    - Logic Programming - Prolog
6. Uncertain Knowledge and Reasoning - Quantifying Uncertainty
    - Acting under Uncertainty
    - Uncertainty and Probability
    - Syntax and Semantics
    - Inference by Enumeration, Normalization
7. Uncertain Knowledge and Reasoning - Probabilistic Reasoning (optional)
    - Bayesian Networks
    - Hidden Markov Models
    - Kalman Filters
8. Learning - Theory of Learning
    - General Learning Model
    - Inductive Learning
    - Learning Decision Trees
    - Artificial Neural Networks (Perceptrons, RBF, Deep Learning)
9. Learning - Unsupervised Learning
    - Clustering, K-Means, EM Algorithm
    - Principal Component Analysis
    - Applications: Compression, Feature Representation
10. Learning - Supervised Learning
    - Bayes Classifier, Bayes Decision Rules
    - Discriminant Analysis
    - Applications: Feature Extraction for Classification
11. Learning -Probabilistic Models
    - Statistical Learning Theory (STL)
    - Structural Risk Minimization (SRM)
    - Support Vector Machines (SVM)
12. Learning - Other Popular Models (optional)
    - Bayesian Learning
    - Genetic Algorithms
    - Reinforcement Learning
13. Perception - Search in Spatial Domain and Frequency Domain (optional)
    - FFT, Lowpass and Highpass Filtering, Convolution Theorem
    - Edge Detection, Line and Curve Detection (Hough Transform)
    - Pictorial Information Search using Geometric or Frequency Features
14. Action - Robotics (optional)
    - Sensors and Vision
    - Path Planning
    - Moving and Control

Grading Policy

- Homework 20%
- Midterm exam 20%
- Project and presentation (topics are related to our course **Contents**) 20%
- Class attendance and participation 10%
- Final exam 30%

Statement on academic integrity:

"Academic Integrity is the cornerstone of higher education and is central to the ideals of this course and the university. Cheating is strictly prohibited and devalues the degree that you are working on. As a member of the NJIT community, it is your responsibility to protect your educational investment by knowing and following the academic code of integrity policy that is found at:

http://www5.njit.edu/policies/sites/policies/files/academic-integrity-code.pdf.

Please note that it is my professional obligation and responsibility to report any academic misconduct to the Dean of Students Office. ***Any student found in violation of the code by cheating, plagiarizing or using any online software inappropriately will result in disciplinary action. This may include a failing grade of F, and/or suspension or dismissal from the university.*** *If you have any questions about the code of Academic Integrity, please contact the Dean of Students Office at dos@njit.edu"*

The "Best Practices" document developed and published on the Provost's website (on the policies page) or directly at http://www5.njit.edu/provost/sites/provost/files/lcms/docs/Best_Practices_related_to_Academic_Integrity.pdf.

Miscellaneous

- Berkeley AI Course Materials
- Prolog:
  - J.R. Fisher, *The Prolog Tutorial*
- Lisp:
  - Paul Graham, *ANSI Common Lisp*, Prentice Hall, 1995.