

Core Courses Syllabi

AI701 - Artificial Intelligence

Title	Artificial Intelligence
Code	AI701
Loading	4 Credit-hours
Prerequisites	<ul style="list-style-type: none"> • Understanding of concepts in Discrete Mathematics, and Probability and Statistics • Proficiency in Python and Pytorch or equivalent library
Catalog Description	This course provides the students a comprehensive introduction to modern artificial intelligence (AI), and some of its representative applications. The students will be familiarized with both the historical and recent AI techniques that have proven successful in building practical systems.
Goal	The aim of this course is to provide students with a comprehensive understanding of the modern development of AI foundations and techniques. Students will be able to develop advanced skills to build AI-based solutions for practical problems.
Contents	This course covers the following major modules: (I) machine learning (II) solving problems by searching (III) knowledge representation and reasoning (IV) Markov decision processes (V) Bayesian networks, and (VI) AI applications
Recommended Textbooks	S. Russell and P. Norvig, <i>Artificial Intelligence: A Modern Approach</i> , 3rd Edition, Prentice Hall, 2010.
Recommended References & Supplemental Material	<p>Relevant research papers, tech reports, and surveys for each topic, where needed, are identified in the teaching plan ahead. In addition, the following textbook may be useful:</p> <p>P. C. Jackson Jr., <i>Introduction to Artificial Intelligence</i>, 3rd Edition, Dover Publications, 2019.</p>



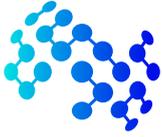
Teaching Week	Topics
1	Course Introduction Lecture <ul style="list-style-type: none">• Course overview (content, learning outcomes, evaluation methods, etc.)• The foundations & history of AI• Agents and environments• The structure of agents Lab <ul style="list-style-type: none">• A quick refresher on Python• Review/discuss real-life AI systems from the agent-environment perspective
2	Machine Learning - Part I Lecture <ul style="list-style-type: none">• Supervised learning• Linear models• Support vector machines (SVMs)• Artificial neural networks (ANNs) Lab <ul style="list-style-type: none">• Practice using Decision Trees, ANNs, and SVMs to solve supervised learning problems on toy datasets by filling in the given code template
3	Machine Learning - Part II Lecture <ul style="list-style-type: none">• Generalization• Unsupervised learning• K-means Lab <ul style="list-style-type: none">• Instructor-led demonstration of topics discussed in the week
4	Solving Problems by Searching - Part I Lecture <ul style="list-style-type: none">• Concept of computational problems• Searching as a problem-solving methodology• Uninformed search strategies Lab <ul style="list-style-type: none">• Solve practical problems with uninformed search, A* search, and hill-climbing



Teaching Week	Topics
5	Solving Problems by Searching - Part II Lecture <ul style="list-style-type: none">• Heuristic search strategies• Local search and optimization problems• Introduction to constraint satisfaction problems (CSPs) Lab <ul style="list-style-type: none">• Review/discuss real-life problems from the perspective of games and CSPs• Assignment # 1 due
6	Knowledge Representation and Reasoning - Part I Lecture <ul style="list-style-type: none">• Syntax vs semantics• Horn clauses• Propositional logic & theorem proving• First-order logic Lab <ul style="list-style-type: none">• Describe and discuss real-life scenarios using propositional logic and first-order logic• Design and discuss knowledge-based agents for practical scenarios
7	Knowledge Representation and Reasoning - Part II Lecture <ul style="list-style-type: none">• Introduction to propositional inference and first-order inference• Unification and lifting• Forward chaining• Backward chaining• Resolution Lab <ul style="list-style-type: none">• Design propositional logic and first-order logic representations of given practical problems
8	Game Playing and Adversarial Search Lecture <ul style="list-style-type: none">• Minimax search• Expectimax• Alpha-beta pruning• Game theory Lab <ul style="list-style-type: none">• Assignment # 2 due



Teaching Week	Topics
9	Uncertain Knowledge and Reasoning Lecture <ul style="list-style-type: none">• Introduction to uncertainty & recap on probability• Representing knowledge in an uncertain domain• Bayesian networks• Exact and approximate inference in Bayesian networks• Relational and first-order logic probability models Lab <ul style="list-style-type: none">• Build and discuss Bayesian networks from real-life scenarios• Practice exact and approximate inference in Bayesian networks by writing functions in the given code skeleton
10	Markov Decision Processes - Part I Lecture <ul style="list-style-type: none">• Introduction to Markov decision processes (MDP)• Policy evaluation• Value iteration• Monte Carlo methods Lab <ul style="list-style-type: none">• Model a practical decision making problem as an MDP
11	Markov Decision Processes - Part II Lecture <ul style="list-style-type: none">• Bootstrapping methods• Introduction to reinforcement learning (RL)• Passive RL• Active RL Lab <ul style="list-style-type: none">• Practice exercises on reward design, optimality in RL, and policy iteration
12	Bayesian Networks Lecture <ul style="list-style-type: none">• Bayesian inference• Hidden Markov models• Gibbs sampling• Particle filtering• Learning Bayesian networks• Laplace smoothing Lab <ul style="list-style-type: none">• Implement object tracking using particle filtering



Teaching Week	Topics
13	AI Applications: Natural Language Processing Lecture <ul style="list-style-type: none">• Language models• Text classification• Information retrieval and extraction• Phrase structure grammars• Syntactic analysis and semantic interpretation• Machine translation• Speech recognition Lab <ul style="list-style-type: none">• Implement a text classification system in the given code skeleton
14	AI Applications: Robotics Lecture <ul style="list-style-type: none">• Robotics hardware: sensors and effectors• Robotic perception• Planning certain and uncertain moves• Dynamics and control: reactive control and RL control• Robotic software architectures Lab <ul style="list-style-type: none">• Assignment # 4 due
15	Guest Lecture & Review Lecture <ul style="list-style-type: none">• Guest Lecture on Trending Topic Lab <ul style="list-style-type: none">• Review and Exam Preparation