Artificial Intelligence

Course Rationale

Artificial intelligence (AI) develops programmed agents (systems) that match or outperform people's abilities to make decisions, to learn, and to plan. To do so, AI develops algorithms and methodologies that sense a system's environment, decide what to do given that data, and effect its chosen actions in its environment.

Course Description

This is a graduate-level course on artificial intelligence. It emphasizes fast and clever search heuristics, thoughtful ways to represent knowledge, and incisive techniques that support rational decision making. Application areas will include game playing, natural language processing, and robotics.

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.

Topic List

Topics include but are not limited to:

- Introduction: foundation definitions, classic AI problems, and their solutions, knowledge representation
- Agents and Problems
- State-space search: uninformed search, informed search, local search
- Constraint satisfaction: principles and practices
- Logics for Agents: probabilistic and logical reasoning
- Knowledge Representation
- Constraint satisfaction: principles and practices

- Machine learning: foundation definitions, computational learning theory, major paradigms
- Planning
- Markov Models
- Introduction to more advanced topics (e.g., embodied cognition, cognitive architectures, autonomy)

Learning Objectives

Students who successfully complete this course will be able to:

- Discuss the agent paradigm as the goal of an intelligent machine
- Describe state space search as a mechanism for problem solving, including optimal solutions and their complexity
- Explain the role of caching, reactivity, heuristics, and planning in state space search.
- Define machine learning and describe the specifics of several prominent machine-learning methods (e.g., SVMs, decision trees, Bayes nets, artificial neural networks, genetic algorithms)
- Evaluate the complexity of an approach to a specific problem and its realistic impact.
- Describe and illustrate the role of constraint satisfaction in AI, with appropriate examples.
- Discuss the role of probabilistic reasoning and mechanisms that employ it
- Discuss the role of logical reasoning and mechanisms that support it

The course has a general goal providing a capability for Empirical AI research that 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.

Assessment

- Class participation 10% will asses both applied and theoretical topics
- Assignments 60% will assess the theoretical topics
- Term project (presentation and report) 30% will assess the applied topics