

COMP 3190 – Introduction to Artificial Intelligence

Course Description

Calendar entry

Principles of artificial intelligence: problem solving, knowledge representation and manipulation; the application of these principles to the solution of 'hard' problems.

Prerequisite: one of COMP 2140 or COMP 2061.

General Course Description

In the 1950s, researchers in the then-new field of Computer Science identified five domains of human intelligence where computer systems could be applied: knowledge, reasoning, linguistic communication, planning, and learning and adaptivity. In the years since, many different strategies have been applied to solve problems in these domains, from logical and symbolic to statistical and connectionist.

This course will introduce the field of Artificial Intelligence by examining the principles underlying symbolic and subsymbolic AI (representation, search, and connectionism), surveying a variety of problems across these domains, and examining more deeply some specific solutions to those problems using several techniques.

Detailed Prerequisites

Before entering this course, a student should be able to:

- Design, implement, and build software requiring multiple modules in an OOP or procedural programming language.
- Design and implement iterative and recursive algorithms, using greedy and divide-and-conquer strategies.
- Implement and apply common abstract data types such as stacks, queues, priority queues, and graphs.
- Implement common data structures and their associated algorithms, such as list operations and tree and graph traversals.
- Compare algorithms according to their worst-case running time analysis and recognize how that will affect their ability to scale with increasing data set sizes.
- Recognize and apply basic statistical concepts such as probabilities.

Course Goals

By the end of this course students will:

- Identify the differences between symbolic and subsymbolic approaches to AI.
- Recognize the current state-of-the-art in solutions and limitations for AI applications.
- Implement strategies for solving AI problems symbolically, like puzzle solving, game playing, and planning.
- Build reasoning systems using logic programming, rules, and uncertainty.
- Identify the components of problems in some example applications of AI.
- Describe the general types and strategies of machine learning systems.

Learning Outcomes

Symbolic AI

Students should be able to:

1. Identify the types of problems that can be represented symbolically.
2. Represent a problem using states and operations that transform between states.
3. Construct a state space from an initial problem state.
4. Use uninformed, informed, heuristic, and adversarial search strategies to find a goal in a state space.

Knowledge Representation and Knowledge-Based Systems

Students should be able to:

1. Represent facts and relationships in a domain using first-order logic.
2. Solve problems in a domain using first-order logic and inference.
3. Implement reasoning systems to resolve queries in a logic programming language.
4. Describe how a logic programming language resolves queries.
5. Express high-level knowledge in a narrow domain using rules.
6. Construct a knowledge-based system to solve problems using rules.

Uncertainty

Students should be able to:

1. Recognize different strategies for reasoning using incomplete knowledge.
2. Incorporate uncertainty into a reasoning system.
3. Represent uncertainty using probabilities.

4. Use probabilistic inference and Bayes' Rule to solve problems with uncertainty.

Application: Planning

Students should be able to:

1. Represent the state of a planning problem as a structured collection of properties.
2. Implement solutions to planning problems using strategies learned earlier in the course.

Application: Natural Language Processing

Students should be able to:

1. Identify the different tasks of NLP systems (understanding, generation, translation) and common problems that occur within each.
2. Identify the components of the study of natural language, and the challenges that occur in the analysis of each, and as a whole.
3. Describe the relationship between semantic analysis and knowledge representation.

Introduction to Machine Learning

Students should be able to:

1. Describe the differences between symbolic and subsymbolic representations.
2. Recognize how learning can be used in a symbolic AI system.
3. Identify the architecture and processing of a simple connectionist system, such as a feedforward neural network.
4. Describe the differences between supervised, unsupervised, and reinforcement learning.