THOMPSON RIVERS UNIVERSITY

Course Outline

Department of Computing Science
Faculty of Science

COMP 3710 - 3
Applied Artificial Intelligence (3,1,0)
Fall 2015

Instructor: Phone/Voice Mail:
Office: E-Mail:

Course Description :

Students investigate non-deterministic computer algorithms that are used in wide application areas but cannot be written in pseudo programming languages. Non-deterministic algorithms have been known as topics of machine learning or artificial intelligence. Students are introduced to the use of classical artificial intelligence techniques and soft computing techniques. Classical artificial intelligence techniques include knowledge representation, heuristic algorithms, rule based systems, and probabilistic reasoning. Soft computing techniques include fuzzy systems, neural networks, and genetic algorithms.

Educational Objectives/Outcomes

Upon successful completion of the course, the student will demonstrate the ability to:

1. Understand the major areas and challenges of AI.
2. Identify problems that are amenable to solution by AI methods, and which AI methods may be suited to solving a given problem.
3. Formalize a given problem in the language/framework of different AI methods.
4. Implement basic AI algorithms.
5. Apply basic AI knowledge and algorithms to solve problems.
6. Design simple software to experiment with various AI concepts and analyse results.

Prerequisites

- COMP 2230 Data Structures, Algorithm Analysis, and Program Design
- STAT 2000 Introduction to Statistics

Required Texts/Materials

- Artificial Intelligence Illuminated, Ben Coppin, Jones and Bartlett Illuminated Series.

Other Available/Recommended Resources

- Norvig P. Russell S., Artificial Intelligence: A Modern Approach, Prentice Hall.

Syllabus – Lecture Topics:

(Note: Not necessarily in this exact order and duration, but very close)

- Part I – Introduction to Artificial Intelligence 1 weeks
  - A Brief History of Artificial Intelligence — Chapter 1
  - Uses and Limitations — Chapter 2
  - Problem characteristics — Lecture note
  - Nature of agents — Lecture note
- Part II – Classical Artificial Intelligence 5 weeks
  - Knowledge Representation — Chapter 3
  - Searching — Chapter 4, 5, 14
    - Search Methodologies
    - Advanced Search
    - Genetic Algorithms (not much classical)
  - Knowledge Representation and Automated Reasoning — Chapter 7, 8, 9
    - Propositional and Predicate Logic
    - Inference and Resolution for Problem Solving
    - Rules and Expert Systems
- Part III – Machine Learning 4 weeks
  - Introduction — Chapter 10
  - Neural Networks — Chapter 11
  - Probabilistic Reasoning — Chapter 12
  - Artificial Life — Chapter 13
- Part IV – Advanced Topics 3 weeks
  - Fuzzy Reasoning — Chapter 18, Lecture note
  - Intelligent Agents — Chapter 19

- Introduction to Understanding Language (when time permits) Chapter 20
- Introduction to Machine Vision (when time permits) Chapter 21

Syllabus – Seminar/Lab Topics :

- Solving the problems regarding to the concept of artificial intelligence
- Solving problems using A* and advanced heuristics
- Solving a problem using a generic algorithm
- Solving the problems regarding to formal languages
- Solving problems using backward chaining and forward chaining
- Implementation of a fuzzy control system
- Solving a problem using a decision tree
- Solving problems using neural networks
- Solving problems using probabilistic reasoning

ACM / IEEE Knowledge Area Coverage

Knowledge Areas that contain topics and learning outcomes covered in the course

Knowledge Area (core)	**Total Hours of Coverage**
IS-Intelligent Systems	Total 10
IS/Fundamental Issues	1
IS/Basic Search Strategies	4
IS/Basic Knowledge Representation and Reasoning	3
IS/Basic Machine Learning	2
DS-Discrete Structures	Total 4
DS/Basic Logic	4
Knowledge Area (elective)	**Total Hours of Coverage**
IS-Intelligent Systems	Total 3
IS/Advanced Search	1
IS/Advanced Representation and Reasoning	0.5
IS/Reasoning Under Uncertainty	1
IS/Advanced Machine Learning	0.5

Body of Knowledge coverage

KA	Knowledge Unit	Topics Covered	Hours
IS	Fundamental Issues (Core-Tier2)	Overview of AI problems, examples of successful recent AI applications	1
		What is intelligent behavior? • The Turing test • Rational versus non-rational reasoning	
		Problem characteristics • Fully versus partially observable • Single versus multi-agent • Deterministic versus stochastic • Static versus dynamic • Discrete versus continuous	
		Nature of agents • Autonomous versus semi-autonomous • Reflexive, goal-based, and utility-based • The importance of perception and environmental interactions	
IS	Basic Search Strategies (Core-Tier2)	Problem spaces, problem solving by search	4
		Factored representation (factoring state in variables)	
		Uninformed search (breadth-first, depth-first with interactive deepening)	
		Heuristics and informed search (hill-climbing, generic best-first, A*)	
		Space and time efficiency of search	
		Two-player games (introduction to minimax search)	
		Constraint satisfaction (backtracking and local search methods)	
IS	Basic Knowledge Representation and Reasoning (Core-Tier2)	Review of propositional and predicate logic	3
		Resolution and theorem proving (propositional logic only)	
		Forward chaining, backward chaining	

		Review of probabilistic reasoning, Bayes theorem	
IS	Basic Machine Learning (Core-Tier2)	Definition and examples of broad variety of machine learning tasks, including classification	2
		Inductive learning	
		Simple statistical-based learning, such as Naïve Bayesian Classifier, decision trees	
		The over-fitting problem	
		Measuring classifier accuracy	
IS	Advanced Search (Elective)	Stochastic search – Simulated annealing; Genetic algorithm	
		Implementation of A* search	
IS	Advanced Representation and Reasoning (Elective)	Rule-based Expert systems	
IS	Reasoning Under Uncertainty (Elective)	Review of basic probability	
		Random variables and probability distributions • Probabilistic inference • Bayes' rule	
		Knowledge representations • Bayesian Networks	
IS	Advanced Machine Learning (Elective)	Definition and examples of broad variety of machine learning tasks	
		Nearest-neighbor algorithms	
DS	Basic Logic (Core-Tier1)	Review of propositional and predicate logic	4
		Normal forms (conjunctive and disjunctive)	
		Validity of well-formed formula	
		Propositional inference rules (concepts of modus ponens and modus tollens)	
		Predicate logic – universal and existential quantification	
		Limitations of propositional and predicate logic (e.g., expressiveness issues)	