University of Saskatchewan
Department of Computer Science
176 Thorvaldson Building
110 Science Place, Saskatoon, SK, S7N 5C9, Canada
Telephine: (306) 966-4886, Facimile: (306) 966-4884

Course Syllabus
CMPT 317: Introduction to Artificial Intelligence

Catalogue Description

A survey of Artificial Intelligence techniques and underlying theory. Topics include problem solving and planning, knowledge representation techniques, reasoning engines and expert systems, and a tour of various application areas of Artificial Intelligence including machine learning, natural language processing and high-level computer vision.

Prerequisite(s):	CMPT 260 and 280.
Class Time and Location:	MWF: 2:30pm-3:20pm, Thorvaldson 205A
Website:	Moodle

Instructor Information

Instructor:	Michael Horsch
Contact:	Email: horsch@cs.usask.ca
	Office Phone: 2161
Office:	Thorvaldson 280.5
Consulting Hours:	Hours: TBA
	Location: Spinks Lab

Course Objectives

Artificial Intelligence is the name of a collection of techniques inspired by the goal of understanding and implementing intelligent behaviour. The course will present some of these techniques so that students can apply them when it makes sense to use them, and understand how they will perform in such applications. The phrase “Artificial Intelligence" also evokes connotations and implications about technology, sometimes fictional and exaggerated. We’ll spend a little time discussing these conceptions and misconceptions.

By the end of the course, students will be able to:

- Describe introductory techniques in Artificial Intelligence, including:
  - Heuristic search and adversarial search
  - Logic for knowledge representation and reasoning
  - Reasoning under uncertainty
  - Machine learning
- Apply introductory techniques in Artificial Intelligence (see above) to solve more-or-less realistic problems.
- To demonstrate critical analysis of potential applications of the introductory techniques

A list of topics is found in the Lecture Schedule (below). A list of formal learning outcomes appears on the final page.

Textbook Information

Required Text

The following book (known internationally as "AIMA") is an extraordinarily good text. Do not think your money is wasted buying a copy of it. You can always resell it at the end of term, as it will be used in future offerings. Don't settle for the first edition, which is much inferior.

- Stuart Russell and Peter Norvig. Artificial Intelligence: A Modern Approach. 3rd Edition. Prentice Hall, 2010, 1132 pages.

Student Evaluation

Grading Scheme

Assignments (4 @ 15%)	60%
Midterm Exam (week of Mar 9)	10%
Final Exam	30%
Total	100%

Criteria That Must Be Met To Pass

The grade will be calculated as above; no other criteria will be used. The assignments carry the bulk of your grade, and will exercise your ability to apply the course material to problems; students should aim for 50–60 marks out of 60 in total on the assignments. They won't be deliberately hard, so reasonable and professional effort should earn grades in that range. The midterm and final will assess the degree to which students can demonstrate critical analysis of the course material. In other words, the examination portion of the course will not test comprehension (that is assumed) but critical analysis.

Attendence Expectation

The expectation is that students behave in a professional manner, including attendance in all lectures, writing the midterm and the final, and submitting every assignment. Exceptions and excuses will be accepted reasonably, and attendance will not be counted for grades. Treat this course as a professional would.

Final Exam Scheduling

The Registrar schedules all final examinations, including deferred and supplemental examinations. Students are advised not to make travel arrangements for the exam period until the official exam schedule has been posted.

Note: All students must be properly registered in order to attend lectures and receive credit for this course.

Course Overview

Lectures

The course material will be covered in a combination lecture style content and problem-driven content. Students will be expected to do a little reading to prepare for classes.

Each unit will start with a real-world problem for motivation. Class-time will be spent in constructive analysis and problem-solving, guided by the course material. By the time we've understood the problem, we'll be receptive to the approaches and techniques accumulated in the textbook and elsewhere. Lecture notes may or may not be posted for any given topic; I expect you'll be in class, and you'll have the textbook.

Tutorials

Tutorials will be held weekly in Thorvaldson 205A from 12:30pm-1:20pm, starting January 12. These sessions are not optional.

Midterm Exam

The midterm will be a dry-run of the final exam, but held early so you know what to expect. The exact timing will be negotiated, within the week of March 9.

Assignments

There will be 4 assignments, easy worth 15%. Approximate due dates: Jan 29, Feb 26, Mar 18, Apr 7. The due date for the last assignment is a hard deadline.

Lecture Schedule

The course will cover approximately Chapters 1 through 18 of AIMA, though we will skip over some chapters completely. Other related topics may be inserted into the schedule as time and interest permit.

Topic/Week/Date	Topic/Subtopics
Introduction (Chap. 1 and 2)	• What is Intelligence? • Definitions of Intelligent Agents
Single-Agent Search (Chap. 3 and 4)	• Breadth-first, Depth-first and Iterative Deepening Search • Heuristic Search (A* search) • Stochastic Local Search (Simulated Annealing, Genetic algorithms)
Adversarial Search (Chap. 5)	• Minimax Search, Alpha-beta pruning • Stochastic Games and Expectiminimax • Partially Observable Games and Perfect Information Monte Carlo Search
Constraint Satisfaction (Chap. 6)	• CSP Domain Definitions • Constraint Propagation, Arc Consistency • Constraint Heuristics
Knowledge Representation and Logic (Chap. 7, 8, 9)	• Propositional Logic, Propositional Inference • First-Order Logic, Propositional Inference (Forward chaining, Backward chaining)
Reasoning Under Uncertainty (Chap. 13 and 14)	• Probability • Bayes Rule • Bayesian Networks, Bayesian Inference

Machine Learning (Chap. 18)	• Definition and examples of broad variety of machine learning tasks, including classification • Inductive learning • Simple statistical-based learning, such as Naive Bayesian Classifier, decision trees • The over-fitting problem • Measuring classifier accuracy

Policies

Recording of Lectures

Lectures are being captured and made available on the Echo 360 site accessible through Moodle. You may also use recording devices if you think that will help. The lecture material will be sourced from AIMA and other references which will be cited appropriately.

Late Assignments

Assignment due dates are for the convenience of the marker. There will be a 1 hour grace period after the due date for last minute preparations. After the grace period, 15% of your grade will be deducted, increasing by 15% every day thereafter. This is not punishment for being late; it is motivation for you to hand something in sooner than later.

Reasonable excuses for late assignments will be accepted; deadlines can be negotiated in class to avoid conflict with deadlines for other courses. Practice professional behaviour, and guard your professional reputation.

Missed Assignments

The intent of assignments is to exercise the course objectives in a professional manner. The assignments will require measured and reasonable amounts of effort, with reasonable deadlines. I expect you to hand something in for every assignment, and for that something to demonstrate professional attitudes towards your work. There is no penalty for missing an assignment (apart from a zero grade on that assignment), but you should consider it a professional failure if you miss an assignment.

Missed Examinations

1. "Students who have missed an exam or assignment must contact their instructor as soon as possible. Arrangements to make up the exam may be arranged with the instructor. Missed exams throughout the year are left up to the discretion of the instructor if a student may make up the exam or write at a different time. If a student knows prior to the exam that she/he will not be able to attend, they should let the instructor know before the exam."

2. "Final exams - a student who is absent from a final examination through no fault of his or her own, for medical or other valid reasons, may apply to the College of Arts and Science Dean's office. The application must be made within three days of the missed examination along with supporting documentary evidence. Deferred exams are written during the February mid-term break for Term 1 courses and in early June for Term 2 and full year courses."

   (2007/08. http://www.arts.usask.ca/students/transition/tips.php)

Incomplete Course Work and Final Grades

When a student has not completed the required course work, which includes any assignment or examination including the final examination, by the time of submission of the final grades, they may be granted an extension

to permit completion of an assignment, or granted a deferred examination in the case of absence from a final examination. Extensions for the completion of assignments must be approved by the Department Head, or Dean in non-departmentalized Colleges, and may exceed thirty days only in unusual circumstances. The student must apply to the instructor for such an extension and furnish satisfactory reasons for the deficiency. Deferred final examinations are granted as per College policy.

In the interim, the instructor will submit a computed percentile grade for the course which factors in the incomplete course work as a zero, along with a grade comment of INF (Incomplete Failure) if a failing grade. In the case where the instructor has indicated in the course outline that failure to complete the required course work will result in failure in the course, and the student has a computed passing percentile grade, a final grade of 49% will be submitted along with a grade comment of INF (Incomplete Failure).

If an extension is granted and the required assignment is submitted within the allotted time, or if a deferred examination is granted and written in the case of absence from the final examination, the instructor will submit a revised computed final percentage grade. The grade change will replace the previous grade and any grade comment of INF (Incomplete Failure) will be removed. For provisions governing examinations and grading, students are referred to the University Council Regulations on Examinations section of the Calendar.

(2011 University of Saskatchewan Calendar/Academic Courses Policy)

Academic Honesty

The University of Saskatchewan is committed to the highest standards of academic integrity and honesty. Students are expected to be familiar with these standards regarding academic honesty and to uphold the policies of the University in this respect. Students are particularly urged to familiarize themselves with the provisions of the Student Conduct & Appeals section of the University Secretary Website and avoid any behaviour that could potentially result in suspicions of cheating, plagiarism, misrepresentation of facts and/or participation in an offence. Academic dishonesty is a serious offence and can result in suspension or expulsion from the University.

All students should read and be familiar with the Regulations on Academic Student Misconduct, http://www.usask.ca/secretariat/student-conduct-appeals/StudentAcademicMisconduct.pdf, as well as the Standard of Student Conduct in Non-Academic Matters and Procedures for Resolution of Complaints and Appeals, http://www.usask.ca/secretariat/student-conduct-appeals/StudentNon-AcademicMisconduct.pdf.

Academic honesty is also defined and described in the Department of Computer Science Statement on Academic Honesty: http://www.cs.usask.ca/undergrad/honesty.php.

For more information on what academic integrity means for students see the Student Conduct & Appeals section of the University Secretary Website at: http://www.usask.ca/university_secretary/pdf/dishonesty_info_sheet.pdf

Examinations with Disability Services for Students (DSS)

Students who have disabilities (learning, medical, physical, or mental health) are strongly encouraged to register with Disability Services for Students (DSS) if they have not already done so. Students who suspect they may have disabilities should contact DSS for advice and referrals. In order to access DSS programs and supports, students must follow DSS policy and procedures. For more information, check http://www.students.usask.ca/disability/, or contact DSS at 966-7273 or `dss@usask.ca`.

Students registered with DSS may request alternative arrangements for mid-term and final examinations. Students must arrange such accommodations through DSS by the stated deadlines. Instructors shall provide the examinations for students who are being accommodated by the deadlines established by DSS.

ACM-2013 Learning Outcomes

This course achieves the following learning outcomes (listed alphabetically) from the ACM-2013 Computer Science Curriculum Guidelines (https://www.acm.org/education/CS2013-final-report.pdf):

- Describe Turing test and the "Chinese Room" thought experiment. [Familiarity, IS-FI-1]
- Differentiate between the concepts of optimal reasoning/behavior and human-like reasoning/behavior. [Familiarity, IS-FI-2]
- Determine the characteristics of a given problem that an intelligent system must solve. [Assessment, IS-FI-3]
- Formulate an efficient problem space for a problem expressed in natural language (e.g., English) in terms of initial and goal states, and operators. [Usage, IS-BSS-1]
- Describe the role of heuristics and describe the trade-offs among completeness, optimality, time complexity, and space complexity. [Familiarity, IS-BSS-2]
- Describe the problem of combinatorial explosion of search space and its consequences. [Familiarity, IS-BSS-3]
- Select and implement an appropriate uninformed search algorithm for a problem, and characterize its time and space complexities. [Usage, IS-BSS-4]
- Select and implement an appropriate informed search algorithm for a problem by designing the necessary heuristic evaluation function. [Usage, IS-BSS-5]
- Evaluate whether a heuristic for a given problem is admissible/can guarantee optimal solution. [Assessment, IS-BSS-6]
- Formulate a problem specified in natural language (e.g., English) as a constraint satisfaction problem and implement it using a chronological backtracking algorithm or stochastic local search. [Usage, IS-BSS-7]
- Compare and contrast basic search issues with game playing issues. [Familiarity, IS-BSS-8]
- Translate a natural language (e.g., English) sentence into predicate logic statement. [Usage, IS-BKR-1]
- Convert a logic statement into clause form. [Usage, IS-BKR-2]
- Apply resolution to a set of logic statements to answer a query. [Usage, IS-BKR-3]
- Design and implement a genetic algorithm solution to a problem. [Usage, IS-AS-1]
- Design and implement a simulated annealing schedule to avoid local minima in a problem. [Usage, IS-AS-2]
- Design and implement A*/beam search to solve a problem. [Usage, IS-AS-3]
- Apply minimax search with alpha-beta pruning to prune search space in a two-player game. [Usage, IS-AS-4]
- Compare and contrast genetic algorithms with classic search techniques. [Assessment, IS-AS-5]
- Compare and contrast various heuristic searches vis-a-vis applicability to a given problem. [Assessment, IS-AS-6]
- Apply Bayes' rule to determine the probability of a hypothesis given evidence. [Usage, IS-RUU-1]
- Explain how conditional independence assertions allow for greater efficiency of probabilistic systems. [Assessment, IS-RUU-2]
- Identify examples of knowledge representations for reasoning under uncertainty. [Familiarity, IS-RUU-3]
- List the differences among the three main styles of learning: supervised, reinforcement, and unsupervised. [Familiarity, IS-BML-1]
- Identify examples of classification tasks, including the available input features and output to be predicted. [Familiarity, IS-BML-2]
- Explain the difference between inductive and deductive learning. [Familiarity, IS-BML-3]
- Describe over-fitting in the context of a problem. [Familiarity, IS-BML-4]
- Apply the simple statistical learning algorithm such as Naive Bayesian Classifier to a classification task and measure the classifier's accuracy. [Usage, IS-BML-5]