University of Wisconsin Madison

COMP SCI 540 section 001 Syllabus

Introduction to Artificial Intelligence

COURSE INFORMATION

Introduction to Artificial Intelligence
COMP SCI 540 001 (3 Credits)
Spring 2017-2018 [1184]

Description
Principles of knowledge-based search techniques, automatic deduction, knowledge representation using predicate logic, machine learning, probabilistic reasoning. Applications in tasks such as problem solving, data mining, game playing, natural language understanding, computer vision, speech recognition, and robotics.

Prerequisite(s)
(COMP SCI 300 or 367) and (MATH 211, 217, 221, or 275) or graduate or professional standing or declared in the Capstone Certificate in Computer Sciences for Professionals

Breadths
N - Natural Science

Instruction Mode
Classroom Instruction

Department: COMPUTER SCIENCES
College: Letters and Science

Canvas Course URL

https://canvas.wisc.edu/

2017-2018 Spring [1184]
Term Start Date: Tuesday, 23-Jan-2018 **Term End Date:** Friday, 18-May-2018

Location and Schedule: Noland Zoology Building 132 MWF 11:00 AM-11:50 AM
CRN: 266004280

How the Credit Hours are Met
This class meets for three 50-minute class periods each week over the semester and carries the expectation that students will work on course learning activities (reading, writing, problem sets, studying, etc) for about 2 hours out of classroom for every class period. The syllabus includes additional information about meeting times and expectations for student work.

INSTRUCTORS AND TEACHING ASSISTANTS

Instructor

YINGYU LIANG
YLIANG@CS.WISC.EDU

Instructor Availability

Monday and Friday, 4-5pm, or by appointment

TA Office Hours

TAs:
Shuo Yang, syang439@wisc.edu, TBA
David Merrel, dmerrell2@wisc.edu, TBA
Shiqi Yang, sqyang@cs.wisc.edu, TBA
Ziyun Zeng, ziyun.zeng@wisc.edu, TBA

Peer mentors:
Qiuxuan Wu, qwu79@wisc.edu, TBA
Carter Steffen, csteffen3@wisc.edu, TBA
Weipeng Zhou, wzhou87@wisc.edu, TBA
Rahul Chityala, chityala@wisc.edu, TBA
Utkarsh Jain, utkarsh.jain@wisc.edu, TBA

GRADING AND COURSE MATERIALS

Course Learning Outcomes

1. Uninformed Search Methods – Be able to formulate problem solving tasks as searching, goal test, operators, state-space graph search formulation, closed world assumption, expanding a node, frontier list, partial solution path, solution path, search tree, different search methods, completeness, optimality, admissibility, complexity, detecting repeated states, explored list.
[S4738]

2. Informed Search Methods – Understand heuristic functions, evaluation functions, best-first search, greedy best-first search, beam search, algorithm A, algorithm A*, admissible heuristic, consistent heuristic, better informed heuristic, devising heuristics.
[S4739]

3. Local Search Methods – Local search problem formulation, operators, neighborhood, move set, hill-climbing algorithm, local optima problem, hill-climbing with random restarts, stochastic hill-climbing (simulated annealing), escaping local optima, Boltzman's equation, cooling schedule, genetic algorithms, crossover, mutation, fitness function, proportional fitness selection, population, crowding.
[S4740]

4. Game Playing – Zero-sum games, perfect information games, deterministic vs. stochastic, playing as search, branching, minimax principle and algorithm, static evaluation, alpha-beta pruning, cutoff, best case and worst case of alpha-beta vs. minimax, iterative-deepening, horizon effect, quiescence search, representing non-deterministic games, chance nodes, expectimax value, Monte Carlo tree search.
[S4741]

5. Constraint Satisfaction - Problem formulation in terms of variables, domains and constraints, constraint graph, depth-first search, backtracking with consistency checking, most constrained variable heuristic, most constraining variable heuristic, least constraining value heuristic, min-conflicts heuristic, min-conflicts algorithm, forward checking algorithm, arc consistency algorithm (AC-3).
[S4742]

6. Unsupervised Learning – Inductive learning, unsupervised learning, feature space, feature, attribute, examples, labels, classes, training set, testing set, classification problems, inductive bias, preference bias, hierarchical agglomerative clustering algorithm, single linkage, complete linkage, average linkage, dendrogram, k-means clustering algorithm, cluster center, distortion cluster quality.
[S4743]

7. K-nearest Neighbors and Decision Trees – K-nearest neighbor algorithm, Ockham's razor, decision tree algorithm, information gain, max-gain, entropy, conditional entropy, remainder, overfitting problem, pruning, training set, testing set, tuning set, setting parameters, k-fold cross validation, leave-one-out cross validation, random forests, bagging ensemble learning.
[S4744]

8. Support Vector Machines – Maximum margin, definition of margin, kernel trick, support vectors, slack variables.
[S4745]

9. Neural Networks – Perceptron,LTU, activation functions, bias input, input units, output units, Perceptron learning rule, Perceptron learning algorithm, epoch, weight space, input space, linearly separable, credit assignment problem, multi-layer feed-forward networks, hidden units, sigmoid, ReLU, back-propagation, gradient descent, deep learning, convolutional neural networks, pooling.
[S4746]

10. Reasoning under Uncertainty – Random variable, mutually exclusive, 3 axioms of probability, joint/conditional/prior/posterior probability, full joint probability, degrees of freedom, marginalization, normalization, product rule, chain rule, conditionalized chain rule, Bayes's rule, conditionalized Bayes's rule, addition/conditioning, independence, conditional independence, naïve Bayes.
[S4747]

11. Bayesian Networks – Bayesian network DAG, conditional probability tables, space saving compared to full joint probability distribution table, conditional independence property defined by a Bayesian network, inference by enumeration from a Bayesian network, naïve Bayes classifier as a Bayesian network.
[S4748]

12. Speech Recognition – Phones, phonemes, Bayes's rule, language/acoustic model, bigram/trigram model, Markov assumption, probabilistic finite state machine, Markov model, state transition matrix, π vector, computing conditional probabilities from a Markov model, hidden Markov model, observation likelihood matrix, computing joint/conditional probabilities from an HMM by enumeration.
[S4749]

13. Computer Vision – Viola-Jones face detection algorithm, boosting ensemble learning, AdaBoost algorithm, Eigenfaces algorithm, nearest-neighbor classification, image space, face space, average face, eigenvalues, eigenvectors, dimensionality reduction.
[S4750]

Grading

- Midterm Exam: about 15%
- Final Exam: about 30%
- Homework Assignments: about 55%

Required Textbook, Software, & Other Course Materials

Required text: S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, 3rd ed., Prentice Hall, 2010. Other readings will be made available electronically on the course web page. Programming assignments will be done using the Java programming language.

EXAMS, QUIZZES, PAPERS & OTHER MAJOR GRADED WORK

Exams, Quizzes, Papers & Other Major Graded Work

The midterm exam will cover material in the first half of the course. It will be taken during a 1-hour block in an evening. Students may bring one 8.5" x 11" sheet of notes on both sides, but otherwise the exam is closed-book. A calculator may also be used. Make-up exams must be approved at least one week before the regular exam.

The final exam will cover material in the whole course. It will be taken during a 2-hour block as assigned by the university. Students may bring one 8.5" x 11" sheet of notes on both sides, but otherwise the exam is closed-book. A calculator may also be used. No make-up final exam is possible except as allowed by university policy.

midterm
March 14, Wednesday, 7:15 pm - 8:15 pm, room TBA

final
May 6, Sunday, 10:05AM - 12:05PM, room TBA

Homework & Other Assignments

Homework assignments will consist of written problems and programming problems. Programming problems will require writing code in the Java programming language. All homework is to be completed individually. Students may do their programming assignments using either their own computer or else one of the computers in the Computer Sciences Department's instructional labs. Answers to written problems and Java code that is written by the student will be handed in electronically using the UW-Madison's Canvas system.

Assignment grading questions must be raised with the instructor within one week after it is returned.

Late assignment policy will be posted on the course website.

ACADEMIC POLICIES

ACADEMIC INTEGRITY

By enrolling in this course, each student assumes the responsibilities of an active participant in UW-Madison's community of scholars in which everyone's academic work and behavior are held to the highest academic integrity standards. Academic misconduct compromises the integrity of the university. Cheating, fabrication, plagiarism, unauthorized collaboration, and helping others commit these acts are examples of academic misconduct, which can result in disciplinary action. This includes but is not limited to failure on the assignment/course, disciplinary probation, or suspension. Substantial or repeated cases of misconduct will be forwarded to the Office of Student Conduct & Community Standards for additional review. For more information, refer to https://conduct.students.wisc.edu/academic-integrity/

ACCOMMODATIONS FOR STUDENTS WITH DISABILITIES

McBurney Disability Resource Center syllabus statement: "The University of Wisconsin-Madison supports the right of all enrolled students to a full and equal educational opportunity. The Americans with Disabilities Act (ADA), Wisconsin State Statute (36.12), and UW-Madison policy (Faculty Document 1071) require that students with disabilities be reasonably accommodated in instruction and campus life. Reasonable accommodations for students with disabilities is a shared faculty and student responsibility. Students are expected to inform faculty [me] of their need for instructional accommodations by the end of the third week of the semester, or as soon as possible after a disability has been incurred or recognized. Faculty [I], will work either directly with the student [you] or in coordination with the McBurney Center to identify and provide reasonable instructional accommodations. Disability information, including instructional accommodations as part of a student's educational record, is confidential and protected under FERPA." http://mcburney.wisc.edu/facstaffother/faculty/syllabus.php

DIVERSITY & INCLUSION

Institutional statement on diversity: "Diversity is a source of strength, creativity, and innovation for UW-Madison. We value the contributions of each person and respect the profound ways their identity, culture, background, experience, status, abilities, and opinion enrich the university community. We commit ourselves to the pursuit of excellence in teaching, research, outreach, and diversity as inextricably linked goals.

The University of Wisconsin-Madison fulfills its public mission by creating a welcoming and inclusive community for people from every background – people who as students, faculty, and staff serve Wisconsin and the world." https://diversity.wisc.edu/