CS 182 Artificial Intelligence
Fall 2018

Instructors: Goran Radanovic and Haifeng Xu

MW 3:00-4:15pm, Northwest B103
`https://canvas.harvard.edu/courses/42715`

1 Overview

Artificial Intelligence (AI) is an exciting field that has enabled a wide range of cutting-edge technology, from driverless cars to grandmaster-beating Go programs. The goal of this course is to introduce the ideas and techniques underlying the design of intelligent computer systems. Topics covered in this course are broadly be divided into 1) planning and search algorithms, 2) probabilistic reasoning and representations, and 3) machine learning (although, as you will see, it is impossible to separate these ideas so neatly). Within each area, the course will also present practical AI algorithms being used in the wild and, in some cases, explore the relationship to state-of-the-art techniques. The class will include lectures connecting the models and algorithms we discuss to applications in robotics, computer vision, and related domains.

Specific topics covered include classical graph search methods, constraint satisfaction problems, Markov decision processes, reinforcement learning, robot motion planning, probability theory, Bayes nets, hidden Markov models, filtering, basic optimization, classification, and regression. The course will provide a good foundation for topics covered in advanced AI courses (CS28x). CS 182 complements CS 181, which emphasizes machine learning. Students who take both CS 182 and CS 181, will have a solid background for understanding and contextualizing modern AI research and experience implementing algorithms in several key areas of the field.

Finally, in spite of its practical usefulness this course is also quite fun. AI also has a long history of research into topics like puzzle-solving, game-playing, robotics, and conversational chat-bots. In this spirit, problem sets will include programming intelligent Pac-Man agents and simple simulated robots.

Learning Outcomes Students completing this course should be able to:

- choose the appropriate representation for an AI problem or domain model, and construct domain models in that representation,
- choose the appropriate algorithm for reasoning within an AI problem domain,
- implement and debug core AI algorithms in a clean and structured manner,
- design and analyze the performance of an AI system or component,
- describe AI algorithms and representations and explain their performance,
- and critically read papers on AI systems.

2 Preliminaries

Prerequisites CS 51 (experience with programming and systems building) and familiarity with probability at the level of Stat 110 will be assumed (and may be taken concurrently). Students should also have some exposure to complexity (big-O notation, basic algorithm analysis). No previous exposure to AI is assumed. Talk to the teaching staff or skim the future readings if you are concerned about your preparation. All programming assignments will be in Python.

Textbooks Required textbook, available for purchase at the COOP:

- Russell, Stuart, and Peter Norvig. 2010. *Artificial Intelligence: A Modern Approach*. 3rd ed. Pearson. (Also called AIMA).

During the RL part of the course we will use

- Sutton, Richard S, and Andrew G Barto. 2016. *Reinforcement Learning: An Introduction 2nd Edition (Draft)*. MIT Press.

which is available for free online at `https://web.stanford.edu/class/psych209/Readings/SuttonBartoIPRLBook2ndEd.pdf`.

We will also use

- LaValle, Steven M. 2006. *Planning Algorithms*. Cambridge University Press

which is available for free online at `http://planning.cs.uiuc.edu/`.

Support resources We will be using Piazza for questions. Unless your question would reveal confidential information or give away answers to homework questions, please post there publicly. Otherwise, send us a private message via Piazza. We also strongly encourage you to answer each others questions.

Office hours and email Staff office hours are listed on the course website. You are welcome to come to office hours with specific questions about the material, to discuss final project ideas, or just to chat about things you find interesting and want to explore further. To avoid duplication of questions and keep the email load manageable and centralized, *please use Piazza for all questions*.

Sections There will be two weekly sections (see course website for schedule). Sections will review material covered in the previous weeks lectures (except before midterms, when a more comprehensive review will take place. Attendance at section is strongly encouraged. Section notes will be posted on Canvas.

3 Schedule

Below is a tentative schedule for the semester. Please refer to the updated weekly course schedule posted on the Canvas course website.

CS 182 - Fall 2018 : Course schedule

Date	#	Topic	Reading	Assignments	Section	Cluster
9/5/2018	1	Introduction to AI	AIMA 1, (2)	*P0: Python and Math Review*		Overview
9/10/2018	2	Uninformed Search	AIMA 3.1-3.4		Python and Math Review	Search & Constraint Satisfaction
9/12/2018	3	A* and Heuristics	AIMA 3.5, 3.6	**P0 due 9/14 @ 5 PM**; *P1: Search*		
9/17/2018	4	Adversarial Search and Games	AIMA 5		Search	
9/19/2018	5	Constraint Satisfaction Problems I	AIMA 6.1			
9/24/2018	6	Constraint Satisfaction Problems II	AIMA 6.2-5		Games and CSPs	
9/26/2018	7	Local Search	AIMA 4.1-3	**P1 due @ 11:59 PM**; *P2: Games/CSPs*		
10/1/2018	8	Markov Decision Processes 1	AIMA 17.1-3, S&B 3-4		CSPs and Local Search	Behavior Planning and Control
10/3/2018	9	Markov Decision Processes 2	``			
10/8/2018		**Columbus Day**			MDPs	
10/10/2018	10	Reinforcement Learning 1	AIMA 21, S&B 6.1-3,5	**P2 due @ 11:59 PM**; P3: MDPs/RL		
10/15/2018	11	Reinforcement Learning 2	``		RL	
10/17/2018	12	Robot motion planning 1	LaValle (3),4,5			
10/22/2018	13	Robot motion planning 2	`` Optional: https://arxiv.org/pdf/1105.1186.pdf		Midterm #1 review	
10/24/2018	G1	Guest Lecture: Ethics of AI		**P3 due @ 11:59 PM**; *P4: HMMs/Planning*		
10/29/2018		**Midterm #1 Covering L1-11**			Motion Planning	
10/31/2018	14	Hidden Markov Models (HMMs) 1	AIMA 13.1-5, 15.2-5			Probabilistic Reasoning
11/5/2018	15	HMMs 2 + Bayes Nets 1	AIMA 14.1-2		HMMs	
11/7/2018	16	Bayes Nets 2	AIMA 14.4-5	**P4 due 11/ 9 @ 5 PM**; *P5: Bayes/Classify*		
11/12/2017		**Veterans Day**		**Project Proposal Due @ 11:59 PM**	Bayes Nets	
11/14/2017	17	Classification and regression	AIMA 18.1-2,6, 20.1-2.2	**Turing Paper Response due @ 11:59 PM**		Machine Learning
11/19/2018	G2	Guest Lecture: Classic Turing paper	Turing Paper		N/A	
11/21/2018		**Thanksgiving**				
11/26/2018	18	Clustering	AIMA 18.8	**Project Update Due @ 11:59 PM**	Classify Regress Cluster	
11/28/2018	G3	Gues Lecture: Goal Recognition		**P5 due @ 11:59 PM**		
12/3/2018	19	Deep Learning + Course Wrap-up			Midterm #2 review	
12/5/2018		**Midterm #2 Covering L12-18**				
TBD		Student Poster Presentations	Location TBD	**Poster Presentation @ TBD**		
12/10/2018		Project Reports Due		**Project Reports due @ 11:59 PM**		

4 Course Requirements

The course has several components:

- 40%: Two in-class exams
- 35%: Six problem sets, and 1-3 paper reading responses
- 25%: A final project (done in groups of 2-3)

Final grades take into account each component. You must achieve a passing grade in all components to pass this course. Although we won't publish hard grade cutoffs, just note that to receive an "A" you must have high performance is all categories.

Problem Sets The 6 problem sets (P0–P5) will be published on the course webpage. Most problem sets have two components: programming and written. The programming part can be done in pairs or individually. The written part will focus more on conceptual/analysis questions and must be done individually. Computational assignments will ask you to develop implementations of algorithms discussed in class. We expect that all code will run, be well-written and be commented appropriately. All problem sets will be submitted through Gradescope.

Paper Reading Responses We will have 1-3 paper reading responses. These will be short responses to help guide guest lecture discussions (submitted as a <1 page PDFs to Canvas).

Late days Each student is allotted a total of **five** late days for use on **problem sets only**. They extend the due date by **24 hours** and a **maximum of 2 late days** can be used towards any individual assignment. **Late days may not be used on any part of the final project or paper responses.**

Grading Written assignments will be hand graded. Programming assignments can be submitted and autograded an unlimited amount of times until the due date (or late due date). If you have used up your 5 late days, you will be penalized 25% per day, up to two days max, with no credit after two days. In cases of medical or other emergencies which interfere with your work, have your Resident Dean contact the instructor. Any grading disputes on written assignments must be submitted as Piazza private messages ***within one week of the grades being posted*** after which the grade is final. Except in extraordinary circumstances, ***no regrades will be accepted on programming assignments***. The work will be fully regraded and the grade may go up or down. Grades are based on correctness, performance, depth of analysis, documentation, and clarity.

Readings Each class meeting is preceded by a reading assignment, which will be assumed during the lecture and discussion in class. You should set aside 2 hours to compete each reading. We do not expect you to fully understand everything before coming to class, but the goal is to prepare for class, familiarize yourself with new terminology and definitions, and to determine which part of the subject you want to hear more about. We encourage you to bring questions to class about material that is confusing. Other students might share your confusion.

In-class Exams There are two in-class exams (closed book, no notes), one covering the first half of the course material and the second covering the second half of the course material. See the schedule for dates and topics covered. Reviews will be held preceding each exam.

Final Project During the second half of the course students will design and carry out a final project in groups of 2-3 students. The final project is of your choosing: it can describe a system you have built or discuss more theoretical issues or even survey cutting edge work in an active area of AI research. We will provide a list of potential topics. Most people who have taken the course consider this one of the most fun and rewarding parts of the course, and we hope you'll have fun with it too. Students will have to submit a proposal, an update, and give a poster presentation, which will all allow the teaching staff to provide feedback before submission of the final report at the end of reading period. Attendance at the poster session is mandatory.

Collaboration Policy Each assignment will include a programming component and a written component. P0 must be done individually. The *programming component* of P1–5 can done and submitted in pairs. In pairs implies designing and writing the code together and both submitting the same code and receiving the same grade. Note that we will treat pairs/non-pairs the same from a grading perspective. We expect you and your partner to design and implement the solutions together. You may also consult with your classmates in other groups as you work on the problem, but you should not talk in terms of pseudocode or real code, and you should not share answers. In addition, you must cite any books, articles, websites, lectures, etc. that have helped you with your work. Similarly, you must list the names of students from other groups with whom your group has collaborated. If you are doing the computational assignment individually, then the same rules apply as for the written assignments: talking is ok, sharing code is not.

The written component of all assignments must be done individually, and each person must submit her/his own written assignment. You are encouraged to consult with your classmates as you work on the problems for the written assignments. However, you should not share answers. After discussions with your peers, make sure that you can work through the problem yourself and ensure that any answers you submit for evaluation are the result of your own efforts. Cite any books, articles, websites, lectures, etc. that have helped you with your work and list the names of students with whom you have collaborated. Note that understanding the concepts in the written assignments is important both for the computational components and the exams.

Final projects must be done in groups of 2-3 students. You are encouraged to discuss your project ideas with your peers. For any questions not covered in this document, please contact the course staff via Piazza for clarification.

Diversity and Inclusion In an ideal world, science would be objective. However, much of science is subjective and is historically built on a small subset of privileged voices. We acknowledge that it is possible that there may be both overt and covert biases in the material due to the lens with which it was written, even though the material is primarily of a scientific nature. Since integrating a diverse set of experiences is important for a more comprehensive understanding of science please contact the course staff (in person or electronically) or submit anonymous feedback if you have any suggestions to improve the quality of the course materials.

We would like to create a learning environment that supports diversity of thoughts, perspectives, and experiences, and honors your identities. If you have a name and/or set of pronouns that differ from those that appear in your official records, please let us know! If you feel like your performance in the class is being impacted by your experiences outside of class, please don't hesitate to contact us. If you prefer to speak with someone outside of the course, the SEAS Director of Diversity, Inclusion and Belonging is an excellent resource.