CS 5100: Foundations of Artificial Intelligence (Fall 2019)

1 General Information

Time: Monday, Thursday 11:45–1:25 Location: International Village 019

2 Teaching Staff

- **The preferred platform for asking questions and contacting staff is Piazza.**
- If e-mail contact is necessary (e.g., sending attachments), **the preferred e-mail address that reaches all staff is** `cs5100-staff@ccs.neu.edu`.
- Only e-mail individual staff if absolutely necessary (e.g., confidential issue), and note that response will typically be slower than contacting all staff via Piazza or the staff mailing list.

Role	Name and E-mail	Office Hours	Location
Instructor	Lawson L.S. Wong `lsw@ccs.neu.edu`	Tue 10-12 and by appointment	513 ISEC
TA	Sabbir Ahmad `ahmad.sab@husky.neu.edu`	Thu 2:30-4:30	242 Forsyth Building
TA	Durga Chillakuru `chillakuru.d@husky.neu.edu`	Sat 10:30-12:30	1st Floor, WVH
TA	Shlok Abhay Gandhi `gandhi.shl@husky.neu.edu`	Fri 10-12	302 Kariotis Hall
TA	Akshay Patil `patil.aksh@husky.neu.edu`	Wed 3:30-5:30	225 Richards Hall

3 Course Overview

This course will introduce students to the fundamentals of artificial intelligence, including the broad areas of search, decision-making under uncertainty, graphical models (reasoning under uncertainty), and machine learning.

The above topics only cover a small portion of the entirety of AI, and do not even cover all of the fundamentals (prominent topics that will not be covered include logical reasoning and knowledge representation). However, by the end of the course, students will have developed a sufficiently broad set of technical tools, that will enable them to solve many real-world problems, self-learn additional techniques, and pursue further specialized courses in AI.

The course material will focus on problem types, models, and algorithms. Applications will be discussed when relevant, but will not be the focus of the content. However, in the spirit of experiential learning, there will be significant opportunities for implementation and application, through the programming assignments and the project.

4 Textbook and Reference Materials

There is no required textbook. However, the following materials are recommended:

- The standard AI textbook is *Artificial Intelligence: A Modern Approach* (AIMA; 3rd edition), by Stuart Russell and Peter Norvig. This textbook serves as excellent reference material, and this course builds on the book's content. If you are considering pursuing further studies in AI, obtaining and reading this textbook is highly recommended. The 2nd edition is reasonably similar to the current edition (with different chapter/section numbers).
- This course is heavily influenced by the CS 188 course at UC Berkeley, developed by Dan Klein and Pieter Abbeel. They offer additional lecture slides, videos, and practice problems on similar topics at: `http://ai.berkeley.edu/course_schedule.html`

5 Prerequisites

- All programming assignments must be completed in Python 2.7.
- Later in the course you will need basic probability and single-variable differential calculus. A short refresher will be provided, but it would help to learn this as soon as possible.

6 Announcements and Discussion

Course material and announcements will be posted on Piazza. The site also offers an excellent discussion forum, where both instructors and fellow students can answer questions. Everyone is encouraged to participate. Questions/notes can be posted anonymously or with identity, and may also be posted privately only to instructors. Note that posting questions/notes via Piazza will most likely result in faster responses compared to e-mailing individual instructors. Piazza sign-up link: `http://piazza.com/northeastern/fall2019/cs5100`

Grades will be posted on Blackboard.

7 Coursework

Type	Frequency	Due dates
Exercises	$\sim$ Biweekly (5 total)	Monday (beginning of class)
Programming assignments	$\sim$ Biweekly (4 total)	Monday (beginning of class)
Exam	1 total	November 21
Project	1 total	See schedule below

- Exercises are based on the previous two weeks of material. Students may discuss the problems with other students, but must write up their own solutions. On each assignment, please also indicate who you discussed with (if any).
  Lateness: Up to one day late (24-hour period), penalized by 10%.
- Programming assignments are designed to let you see algorithms working in practice. Students should work on this by themselves. Resist the temptation to search for existing solutions – the process of implementation and debugging is critical to learning the material.
  Lateness: Up to one day late (24-hour period), penalized by 10%.
- The project offers an opportunity to apply learned techniques on a substantial problem that interests the student. Further details and (non-exhaustive) topic suggestions will be provided in October. Here is a rough timeline for the project, but is subject to change:
  - October 24: Project proposal due
  - October 31: Review proposal with TAs
  - November 11: Milestone 1
  - November 25: Milestone 2
  - December 2/5: Presentation
  - December 6: Draft report
  - December 9/10: Interview / debriefing
  - December 13: Final report

8 Academic Integrity

Cheating and other acts of academic dishonesty will be referred to OSCCR (office of student conduct and conflict resolution) and the Khoury College of Computer Sciences.

9 Schedule (subject to change; version 20191001)

Date	#	Topic	Reference (AIMA)	Assignments due
9/5	1	Course overview	Ch. 1–2	
9/9	2	Uninformed search	3.1–3.4	*Course component weights*
9/12	3	Uninformed search (continued);	3.4	
		Heuristic search	3.5	
9/16	4	Heuristic search (continued)	3.5–3.7	Ex 1
9/19	5	Heuristic search (continued);	3.5.2	PA 1 (parts 1–3, 5)
		Adversarial search	5.1–5.2, 5.4	
9/23	6	Adversarial search (continued)	5.5	PA 1 (all parts)
9/26	7	Probability;	13.1–13.3	
		Decision theory	16.1–16.3	
9/30	8	Markov decision processes (MDPs)	17.1–17.2	Ex 2
10/3	9	Solving MDPs	17.2–17.3	
10/7	10	Reinforcement learning	21.1–21.2	PA 2
10/10	11	Reinforcement learning (continued)	21.2–21.4	
10/14		Columbus Day (no class)		Ex 3
10/17	12	Bayesian inference;	13.3–13.5	
		Hidden Markov models (HMMs)	15.1	
10/21	13	Inference in HMMs	15.2–15.2.1, 15.5.3	PA 3
10/24	14	Bayesian networks	14.1–14.2	*Project proposal*
10/28	15	Inference in Bayesian networks	14.4–14.4.2	Ex 4
10/31		*Project proposal reviews*		
11/4	16	Machine learning	18.1–18.2, 18.4, 18.11	PA 4
11/7	17	Linear regression	18.6–18.6.2	
11/11		Veterans' Day (no class)		*Project milestone 1*
11/14	18	Logistic regression	18.6.3–18.6.4, 18.11	
11/18	19	Artificial neural networks	18.7	Ex 5
11/21		**Exam in ISEC Auditorium**		
11/25	20	The really hard problems	Ch. 22–27	*Project milestone 2*
11/28		Thanksgiving Day (no class)		
12/2		*Project presentations*		
12/5		*Project presentations*		*Draft report due 12/6*
				Final report due 12/13

Beginning with the material on Bayesian networks (10/24) and continuing into machine learning, we go into more depth than the textbook. Separate notes will be provided on these topics.

10 Learning Objectives

Module 1: Search – Sequential decision-making under certainty

- Concepts: Agents, environments, states, actions, graph search, tree search, heuristics
  Further types of search problems: Constraint satisfaction, adversarial search (game tree)
  Properties of search algorithms: Time/space complexity, optimality, soundness, completeness

- Algorithms: BFS, DFS, IDS, UCS, Greedy search, A* search
  Understand how the above algorithms are all unified by a priority queue
  Further algorithms: Minimax search, expectimax search

- Mastery objective: Given a sequential decision-making problem (with deterministic outcomes), formulate it as a search problem (by specifying formal components of a search problem), and solve the problem with an appropriate algorithm and heuristics (if applicable).

Module 2: MDPs – Sequential decision-making under uncertainty

- Concepts: Probability (expectation, conditional), utility, maximum expected utility
  Markov decision process (MDP): Reward, return, value, policy, Bellman equation
  Reinforcement learning (RL): Exploration vs. exploitation, model-based vs. model-free, Q-function (action-value function), temporal difference error, linear function approximation

- Algorithms: Value back-up (and-or tree), value iteration (dynamic programming), Q-learning

- Mastery objective: Given a sequential decision-making problem (with stochastic outcomes), formulate it as an MDP (by specifying formal components of an MDP), and solve the problem with an appropriate MDP/RL algorithm.

Module 3: Bayesian networks – Reasoning under probabilistic uncertainty

- Concepts: Probability (conditional, marginalization), Bayes' rule, Bayesian inference
  Hidden Markov models (HMM): Belief, filtering, conditional independence, particle filtering
  Bayesian networks: Representation, conditional independence assumptions, exact inference

- Algorithms: Forward filtering, particle filtering, d-separation, variable elimination

- Mastery objective: Given an inference problem (with unobserved random variables), represent it as a graphical model (Bayesian network / HMM) with suitable conditional independence assumptions, and solve the inference query with an appropriate Bayesian inference algorithm.

Module 4: Learning – Acquisition of knowledge (under uncertainty)

- Concepts: Categories of machine learning, supervised learning, regression, classification
  Learning problem: Dataset, hypothesis, parameters, loss/error function, learning algorithm
  Further concepts: Maximum likelihood, hyperparameters, model selection, train/validate/test

- Models: Linear regression, logistic regression, (deep) neural networks
  Algorithms: Coordinate descent, (stochastic) gradient descent, cross validation

- Mastery objective: Given a learning problem with an optimization objective, derive an appropriate learning algorithm, and understand how to apply the algorithm in practice.