CS 4100: Artificial Intelligence (Fall 2023)

1 General Information

Time: Monday, Wednesday 2:50–4:30 | Location: Snell Engineering Center 108

2 Teaching Staff

- **The preferred platform for asking questions and contacting staff is Piazza.** `https://piazza.com/northeastern/fall2023/cs4100sec1`
- If e-mail contact is necessary (e.g., sending attachments), **the preferred e-mail address that reaches all staff is** `cs4100-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`	Fri 2:30–4:30 PM and by appointment	513 ISEC and Zoom 617 373 7459
TA	Alesia Chernikova `chernikova.a@northeastern.edu`	Thu 5–7 PM	Snell Library 045 and Zoom 628 504 2241
TA	Ketaki Nitin Kolhatkar `kolhatkar.k@northeastern.edu`	Mon 5:30–7:30 PM	Zoom 926 2894 7554
TA	Yu Qi `qi.yu2@northeastern.edu`	Wed 5–7 PM	Snell Library 045 and Zoom 757 826 7655
TA	Oliver Toh `toh.o@northeastern.edu`	Tue 4–6 PM	Zoom 221 470 5816

3 Course Overview

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

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; 4th 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 3rd edition is reasonably similar to the current edition (with different chapter/section numbers).
  `http://aima.cs.berkeley.edu/index.html`
- This course is heavily influenced by the CS 188 course at UC Berkeley, developed by Dan Klein and Pieter Abbeel and many others.
  - The course archives offer additional lecture slides, videos, and practice problems:
    `https://inst.eecs.berkeley.edu/~cs188/archives.html`
  - Lecture videos from CS 188 Summer 2016:
    `https://www.youtube.com/channel/UCHBzJsIcRIVuzzHVYabikTQ/videos`

5 Prerequisites

- All programming assignments must be completed in Python 3.
- Later in the course you will need to use basic probability. A short refresher will be provided, but it would help to learn this as soon as possible.
- Even later in the course you will need to use basic single-variable differential calculus. A refresher will not be provided.

6 What is Where

- Piazza: `https://piazza.com/northeastern/fall2023/cs4100sec1`
  Content: Announcements, lecture slides/notes, assignments, materials, Q&A, discussion.
  Preferred platform for contacting course staff.
  Piazza is our hub; if in doubt, try to look/ask on Piazza first. 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.
- Canvas: `https://northeastern.instructure.com/courses/152636`
  Content: Assignments, grades.
  The main exception to the "everything is on Piazza" rule is that we will use Canvas for assignments and grades. All assignments should be submitted on Canvas, and all grades/feedback will be posted on Canvas. Assignment files may be duplicated on Piazza for convenience, but the version on Canvas supersedes the rest.

7 Coursework

Type	Frequency	Due dates
Exercises	$\sim$ Biweekly (4 total)	Friday (11:59 PM ET)
Programming assignments	$\sim$ Biweekly (4 total)	Friday (11:59 PM ET)
Exam	2 total	October 18, November 20
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). Also see Collaboration Policy below.
  Additionally, there will be a short written exercise on ethics in AI.
  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%.

- There will be 2 midterm exams, administered in-person, during class time (2:50–4:30 PM on Wednesday, October 18 and Monday, November 20). They will cover the first and second modules of the course respectively, and are not cumulative. See the course schedule for details. There is no final exam (but there is a final project).

- The project offers an opportunity to apply learned techniques on a substantial problem that interests the student. Further details, (non-exhaustive) topic suggestions, and examples of projects from previous semesters will be provided in October.
  Here is a rough timeline for the project, but is subject to change:
  - October 20: Project pre-proposal due
  - October 31: Project proposal due
  - December 1: Milestone
  - December 6: Presentation (optional, for selected groups)
  - December 8: Draft report
  - December 15: Final report

8 Academic Integrity and Collaboration Policy

We encourage collaboration and discussion in the class, as long as help is fully acknowledged. Discussion of high-level ideas is generally permitted (except during exams), but submissions should always be your own work (except project, which should be your team's work). For specific collaboration constraints on different parts of the course, see the Coursework section above.

This collaboration policy applies equally to recent AI tools such as ChatGPT and other large language models (LLMs). You are welcome to explore these tools and engage in discussion with them, but you should not be copying answers/code from them. Additionally, you should acknowledge usage of such tools and include your work (interaction) with them.

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 20231001)

Date	Lec #	Topic	Reference (AIMA 4e)	Assignments due (11:59 PM)
9/6	1	Course overview	1–2	
9/11	2	Uninformed search	3.1–3.4	
9/13	3	Uninformed search (continued)	3.4	
9/15				*Course component weights*
				PA 1 (Q1, Q2)
9/18	4	Heuristic search	3.5	
9/20	5	Heuristic search (continued)	3.5–3.6	
9/22				Ex 1
9/25	6	Adversarial search	5.1–5.2	
9/27	7	Adversarial search (continued)	5.3–5.5	
9/29				PA 1 (all parts)
10/2	8	Probability	12.1–12.2	
		Decision theory	16.1–16.3	
10/4	9	Markov decision processes (MDPs)	17.1–17.2	
10/6				Ex 2
10/9		Indigenous Peoples' Day (no class)		
10/11	10	Bellman equation	17.1–17.2	
10/13				PA 2
10/16	11	Value iteration	17.2	
10/18		**Exam 1**	Lectures 2–7	
10/20				*Project pre-proposal*
10/23	12	Reinforcement learning	17.3, 22.1–22.2	
10/25	13	Q-learning	22.2–22.4, 22.7.1	
10/27				Ex 3
10/30	14	Deep Q-networks (DQN)	22.4.3, 22.7.1	
		Machine learning	19.1–19.2	
		Linear regression	19.6–19.6.2	
10/31				*Project proposal*
11/1	15	Logistic regression	19.6.4–19.6.5, 20.2.1	
11/3				PA 3
11/6	16	The practice of machine learning	19.4, 19.9	
11/8	17	Neural networks and deep learning	21.1–21.2, 21.4	
11/10				Ex 4
11/13	18	Ethics 1	27	
11/15	19	Ethics 2	27	
11/17				PA 4
11/20		**Exam 2**	Lectures 8–17	
11/22		Thanksgiving (no class)		
11/27	20	Bayesian inference	12.3–12.5	Ethics
11/29	21	Hidden Markov models (HMMs)	14.1–14.3	
12/1				*Project milestone*
12/4	22	Generative models		
12/6		*Selected project presentations*		
12/8				*Draft report due 12/8*
12/15				*Final report due 12/15*

10 Learning Objectives

Module 1: Search – Sequential decision-making under certainty

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

- 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: Learning – Acquisition of knowledge and inference

- 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
  Inference: Probability (conditional, marginalization), Bayes' rule, Bayesian inference
  Graphical models: Hidden Markov models (HMM), Bayesian networks, generative models

- 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.