CS 4100: Artificial Intelligence (Spring 2021)

1 General Information

Time: Monday, Wednesday 2:50–4:30 Location: Shillman Hall 305

2 Teaching Staff

- **The preferred platform for asking questions and contacting staff is Piazza.** `https://piazza.com/northeastern/spring2021/cs4100`
- 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 1–3 and by appointment	Teams
TA	Shuo Jiang `jiang.shuo@northeastern.edu`	Mon 5–7	Teams
TA	Jung Yeon (John) Park `park.jungy@northeastern.edu`	Thu 12:30–2:30	Teams
TA	Sean Wallace `wallace.se@northeastern.edu`	Sun 3–5	Teams

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; 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`
  - We will use the CS 188 Fall 2020 version of programming assignments:
  `https://inst.eecs.berkeley.edu/~cs188/fa20/projects`

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 Hybrid NUflex

The first semester of Hybrid NUflex was mostly successful; as we continue to explore and develop this new medium of instruction, we would greatly appreciate your patience as things are worked out. The course will be conducted in hybrid format, with the instructor teaching in-person (when possible), and students participating both in-person and remotely. To facilitate this, we will use a combination of tools: Canvas, Teams, and Piazza. (We do not plan on using Zoom.)

- Canvas: `https://northeastern.instructure.com/courses/60150`
  Content: Syllabus, lecture notes, lecture slides, assignments, grades.
  All official content except lectures and office hours will be posted on Canvas. We will use Piazza as the primary platform for announcements, materials, and discussion (see below), but core material (syllabus, notes, slides, assignments) will be duplicated on Canvas. Assignment grades will only be posted on Canvas.

- Teams: `https://teams.northeastern.edu`
  Content: Lecture live-streams and recordings (when possible), office hours.
  All registered students have been added to the "Khoury - CS 4100 1 (Spring 2021)" team. We will live-stream lectures in the "Lectures" channel and hold office hours remotely in the "Office hours" channel. If possible, we would like to record lectures for students in the course to refer to later; we will ask for consent to record at the beginning of each lecture. These recordings will be made available only to students enrolled in the course, instructor of record, and any teaching and instructional assistants assigned to the course.

- Piazza: `https://piazza.com/northeastern/spring2021/cs4100`
  Content: Announcements, materials, Q&A, and discussion.
  Preferred platform for contacting course staff.
  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.

7 Coursework

Type	Frequency	Due dates
Exercises	$\sim$ Biweekly (5 total)	Monday (11:59 PM EST/EDT)
Programming assignments	$\sim$ Biweekly (4 total)	Monday (11:59 PM EST/EDT)
Exam	1 total	March 3
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 midterm exam will be administered remotely during class time (2:50–4:30 PM on Wednesday, March 3). Details will be forthcoming; we expect the exam to be open book, open computer, and possibly even open Internet. No external help / collaboration is permitted.
- 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 February. Here is a rough timeline for the project, but is subject to change:
  - March 15: Project proposal due
  - April 14: Milestone
  - April 21: Presentation (optional, for selected groups)
  - April 23: Draft report
  - April 26/27/28: Interview / debriefing
  - April 30: 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 20210314)

Date	Lec #	Topic	Reference (AIMA 4e)	Assignments due (11:59 PM)
1/20	1	Course overview	1–2	
1/25	2	Uninformed search	3.1–3.4	*Course component weights*
1/27	3	Uninformed search (continued)	3.4	PA 1 (Q1, Q2)
2/1	4	Heuristic search	3.5	Ex 1
2/3	5	Heuristic search (continued)	3.5–3.6	
2/8	6	Adversarial search	5.1–5.2	PA 1 (all parts)
2/10	7	Adversarial search (continued)	5.3–5.4	
2/15		Presidents' Day (no class)		Ex 2
2/17	8	Probability;	12.1–12.2	
		Decision theory	16.1–16.3	
2/22	9	Markov decision processes (MDPs)	17.1–17.2	PA 2
2/24	10	Solving MDPs	17.2	
3/1	11	Solving MDPs (continued);	17.2	
		Reinforcement learning	22.1–22.2	Ex 3
3/3		**Exam** (remote)		
3/8	12	Reinforcement learning (continued)	22.1–22.2	
3/10	13	Reinforcement learning (continued)	22.3–22.4	
3/15	14	Reinforcement learning (continued);	22.4.3, 22.7.1	
		Bayesian inference;	12.3–12.5	
		Hidden Markov models (HMMs)	14.1, 14.3.2	PA 3, *Project proposal*
3/17	15	Inference in HMMs	14.2–14.2.1, 14.5.3	
3/22	16	Bayesian networks	13.1–13.3.2	Ex 4
3/24		Care Day (no class)		
3/29	17	Machine learning	19.1–19.2, 19.4, 19.9	PA 4
3/31	18	Linear regression	19.6–19.6.2	
4/5	19	Logistic regression	19.6.4–19.6.5	Ex 5
4/7	20	Neural networks and deep learning	21.1–21.2, 21.4	
4/12		Care Day (no class)		
4/14	21	Ethics of AI (TBD)	27	*Project milestone*
4/19	22	Ethics of AI (continued) (TBD)	27	
4/21		*Selected project presentations*		
4/23				*Draft report due 4/23*
4/30				*Final report due 4/30*

Beginning with the material on Bayesian networks (3/17) 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
  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: 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.