# CS 4100: Artificial Intelligence (Fall 2021)

## 1 General Information

Time: Tuesday, Friday 1:35–3:15 Location: Shillman Hall 335

## 2 Teaching Staff

- **The preferred platform for asking questions and contacting staff is Piazza.** `https://piazza.com/northeastern/fall2021/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<br>`lsw@ccs.neu.edu` | Fri 3:30–5:30 PM<br>and by appointment | 513 ISEC<br>and Khoury OH |
| TA | Isaac Chan<br>`chan.is@northeastern.edu` | Wed 3–5 PM<br>and Khoury OH | 505 ISEC |
| TA | Shuo Jiang<br>`jiang.shuo@northeastern.edu` | Mon 5–7 PM | Khoury OH |
| TA | Seth Pate<br>`pate.s@northeastern.edu` | Thu 9–11 PM | Khoury OH |

## 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 What is Where

This semester, we return to a "traditional" class format, where most of the course will be conducted in-person. However, office hours will be held virtually via Khoury Office Hours; additionally, the instructor will hold office hours both virtually and physically (see Section 2). To facilitate course logistics, we will use a variety of online tools: Canvas, Piazza, and Khoury OH.

- Piazza: `https://piazza.com/northeastern/fall2021/cs4100`
  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/86735`
  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.

- Khoury Office Hours: `http://www.khouryofficehours.com`
  Content: Office hours.
  Office hours will be conducted mostly virtually, although students are encouraged to attend the professor's office hours in person (Friday 3:30–5:30 PM in 513 ISEC). For virtual OH, we will use the Khoury Office Hours site. You will need a Khoury account to log in. Khoury OH offers a queuing system and may support advanced features such as question grouping in the future. When it is your turn, you will connect to a Microsoft Teams call via the site.

  For more information, see:
  Info: `https://info.khouryofficehours.com`
  Help: `https://info.khouryofficehours.com/help`

  We believe that virtual OH reduces barriers to attending OH. However, if there is sufficient interest/demand for in-person OH, we can consider offering that as circumstances allow.

## 7 Optional Recitation

This course typically does not have recitation sections, but we think that it may be beneficial to have a forum for covering further examples and problems that we cannot fit into regular lecture time. We therefore plan to experiment with a *completely optional* weekly recitation section, where TAs will take turns to cover examples, solve problems, and answer questions. Students who wish to attend are highly encouraged to interact and solve problems *with* the instructing TA. This is not intended as additional office hours, and questions specifically about assignments are discouraged. However, seeing related examples and problems will deepen your understanding.

Time: Wednesday, 1–2 PM
Location: Cahners Hall 008

## 8 Coursework

| Type | Frequency | Due dates |
|---|---|---|
| Exercises | ∼ Biweekly (5 total) | Tuesday/Friday (11:59 PM ET) |
| Programming assignments | ∼ Biweekly (4 total) | Tuesday/Friday (11:59 PM ET) |
| Exam | 2 total | October 19, November 16 |
| 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).
  Additionally, there will be a short written exercise on ethics in AI mid-course.
  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 (1:35–3:15 PM on Tuesday, October 9 and Tuesday, November 16). 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:
  - November 5: Project proposal due
  - December 3: Milestone
  - December 7: Presentation (optional, for selected groups)
  - December 10: Draft report
  - December 13/14/15: Interview / debriefing
  - December 17: Final report

## 9 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.

## 10 Schedule (subject to change; version 20211102)

| Date | Lec # | Topic | Reference (AIMA 4e) | Assignments due (11:59 PM) |
|---|---|---|---|---|
| 9/10 | 1 | Course overview | 1–2 | |
| 9/14 | 2 | Uninformed search | 3.1–3.4 | *Course component weights* |
| 9/17 | 3 | Uninformed search (continued) | 3.4 | PA 1 (Q1, Q2) |
| 9/21 | 4 | Heuristic search | 3.5 | |
| 9/24 | 5 | Heuristic search (continued) | 3.5–3.6 | Ex 1 |
| 9/28 | 6 | Automated planning | 11.1–11.3 | |
| | | Adversarial search | 5.1–5.2 | |
| 10/1 | 7 | Adversarial search (continued) | 5.3–5.5 | PA 1 (all parts) |
| 10/5 | 8 | Probability | 12.1–12.2 | |
| | | Decision theory | 16.1–16.3 | |
| 10/8 | 9 | Markov decision processes (MDPs) | 17.1–17.2 | Ex 2 |
| 10/12 | 10 | Bellman equation | 17.1–17.2 | |
| 10/15 | 11 | Value iteration | 17.2 | PA 2 |
| 10/19 | | **Exam 1** | Lectures 2–7 | |
| 10/22 | 12 | Reinforcement learning | 17.3, 22.1–22.2 | |
| 10/26 | 13 | Q-learning | 22.2–22.4, 22.7.1 | Ex 3 |
| 10/29 | 14 | Ethics 1 | 27 | |
| 11/2 | 15 | Deep Q-networks (DQN) | 22.4.3, 22.7.1 | PA 3 |
| | | Bayesian inference | 12.3–12.5 | |
| 11/5 | 16 | Hidden Markov models (HMMs) | 14.1–14.3 | Ethics, *Project proposal* |
| 11/9 | 17 | Particle filtering | 14.5.3 | |
| | | Bayesian networks | 13.1–13.2 | |
| 11/12 | 18 | Inference in Bayesian networks | 13.2–13.3 | Ex 4 |
| 11/16 | | **Exam 2** | Lectures 8–13 | |
| 11/19 | 19 | Machine learning | 19.1–19.2 | |
| | | Linear regression | 19.6–19.6.2 | |
| 11/23 | 20 | Logistic regression | 19.6.4–19.6.5, 20.2.1 | PA 4 |
| 11/26 | | Thanksgiving (no class) | | |
| 11/30 | 21 | Ethics 2 | 27 | Ex 5 |
| 12/3 | 22 | The practice of machine learning | 19.4, 19.9 | *Project milestone* |
| | | Neural networks and deep learning | 21.1–21.2, 21.4 | |
| 12/7 | | *Selected project presentations* | | |
| 12/10 | | | | *Draft report due 12/10* |
| 12/17 | | | | *Final report due 12/17* |

The optional recitation sections will be held on Wednesdays 1–2 PM, beginning on 9/15.

## 11 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.