# CS 5100: Foundations of Artificial Intelligence (Fall 2018)

## 1 General Information

Time: Tuesday 11:45–1:25, Thursday 2:50–4:30 Location: Robinson Hall 411

## 2 Teaching Staff

| Role | Name and E-mail | Office Hours | Location |
|---|---|---|---|
| Instructor | Lawson L.S. Wong | F 11–1 | 513 ISEC |
| | `lsw@ccs.neu.edu` | | and by appointment |
| TA | Rohan Rajendra Sapre | M 5–7, | Rebecca's Cafe, Churchill Hall |
| | `sapre.r@husky.neu.edu` | W 12–2 | CCIS Lab, WVH (Wed) |
| TA | Bharat Vaidhyanathan | T 2–4, | CCIS Lab, WVH (both T, W) |
| | `vaidhyanathan.b@husky.neu.edu` | W 12–2 | (1st floor common space) |

## 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 and videos on similar topics at: `http://ai.berkeley.edu/course_schedule.html`
- Wikipedia on core AI topics are generally well-written, but may require substitutions in terminology and variable names.
- Visit your teaching staff at office hours! :)

## 5 Prerequisites

- All programming assignments must be completed in Python.
- Later in the course you will need to use basic probability and linear algebra. 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/fall2018/cs5100`

Grades will be posted on Blackboard.

## 7 Coursework

| Type | Frequency | Due dates | Evaluation |
|---|---|---|---|
| Exercises | Biweekly (5 total) | Thursday (beginning of class) | 30% |
| Programming assignments | Biweekly (4 total) | Thursday (beginning of class) | 15% |
| Midterms | 2 total | October 16, November 20 | 30% |
| Project | 1 total | See schedule below | 25% |

- Exercises are based on the previous week's 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 27: Project proposal due
  - October 30/31: Review proposal with TAs
  - November 18: Milestone 1
  - November 29: Milestone 2
  - December 4/6: Presentation
  - December 7: Draft report
  - December 10/11: Interview / debriefing
  - December 14: 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 College of Computer Science.

## 9 Schedule (subject to change; version 20181116)

| Date | Topic | Reference (AIMA) | Assignments due |
|---|---|---|---|
| 9/6 | Course overview; Agents | Ch. 1–2 | |
| 9/11 | Uninformed search | Ch. 3.1–3.4 | |
| 9/13 | Heuristic search | Ch. 3.5–3.7 | |
| 9/18 | Constraint satisfaction | Ch. 6.1–6.3 | |
| 9/20 | Adversarial search (Part 1) | Ch. 5.1–5.2 | Ex 1 |
| 9/25 | Adversarial search (Part 2) | Ch. 5.3–5.5 | PA 1 |
| 9/27 | Probability | Ch. 13.1–13.3 | |
| 10/2 | Decision theory; | Ch. 16.1–16.3 | |
| | Markov decision processes (MDPs) | Ch. 17.1 | |
| 10/4 | Solving MDPs (Part 1) | Ch. 17.1–17.2 | Ex 2 |
| 10/9 | Solving MDPs (Part 2); | Ch. 17.2–17.3 | |
| | Reinforcement learning (Part 1) | Ch. 21.1–21.2 | |
| 10/11 | Reinforcement learning (Part 2) | Ch. 21.2–21.4 | PA 2 |
| 10/16 | Midterm 1 | | |
| 10/18 | Review of Midterm 1, | | |
| | Bayesian inference | Ch. 13.3–13.5 | Ex 3 (due 10/19) |
| 10/23 | Hidden Markov models (HMMs) | Ch. 15.1 | |
| 10/25 | Inference in HMMs | Ch. 15.2–15.2.1 | PA 3 |
| | | | *Project proposal due 10/27* |
| 10/30 | *Project proposal reviews* | | |
| 11/1 | Particle filtering in HMMs; | Ch. 15.5.3 | Ex 4 |
| | Introduction to Bayesian networks | Ch. 14.1 | |
| 11/6 | Bayesian networks | Ch. 14.2 | |
| 11/8 | Inference in Bayesian networks; | Ch. 14.4–14.4.2 | PA 4 |
| | Introduction to machine learning | Ch. 18.1–18.2 | |
| 11/13 | Linear regression | Ch. 18.6–18.6.2 | |
| 11/15 | Logistic regression; | Ch. 18.6.3–18.6.4 | |
| | Training procedures | Ch. 18.4 | Ex 5 |
| | | | *Project milestone 1 due 11/18* |
| 11/20 | Midterm 2 | | |
| 11/22 | Thanksgiving (no class) | | |
| 11/27 | Deep learning and applications | Ch. 18.7, Ch. 22–24 | |
| 11/29 | The really hard problems | Ch. 25–26 | *Project milestone 2* |
| 12/4 | *Project presentations* | | |
| 12/6 | *Project presentations* | | *Draft report due 12/7* |
| | | | *Final report due 12/14* |

Beginning with the material on Bayesian networks (11/1) and continuing into machine learning, we go into more depth than the textbook. Separate notes will be provided on these topics.