# CS188 Syllabus

This syllabus is subject to change. In particular, the midterm date will not be finalized until a week or so into the course.

You may want to look at last semester's slides, but there will be changes.

Note that unreleased project out and due dates are just guesses and will likely change somewhat.

| Day | Topic | Reading | Slides | Video | Assignment | Due |
|---|---|---|---|---|---|---|
| Tu 1/22 | Introduction to AI | Ch. 1 | 2PP · 6PP | video | P0: Tutorial | 1/25 5pm |
| Th 1/24 | Agents and Search | Ch. 3.1-4 (2e: Ch. 3) | 2PP · 6PP | video | | |
| Tu 1/29 | A* Search and Heuristics | Ch. 3.5-6 (2e: Ch. 4.1-2) | 2PP · 6PP | video | HW1 | 2/4 |
| | | | | | P1: Search | 2/8 5pm |
| Th 1/31 | Constraint Satisfaction Problems | Ch. 6.1 (2e: Ch. 5.1) | 2PP · 6PP | video | | |

| Day | Topic | Reading | Slides | Video | Assignment | Due |
|---|---|---|---|---|---|---|
| Tu 2/5 | CSPs II | Ch. 6.2-5 (2e: Ch. 5.2-4) | 2PP · 6PP | video | HW2 | 2/11 |
| Th 2/7 | Game Trees: Minimax | Ch. 5.2-5 (2e: Ch. 6.2-5) | 2PP · 6PP | video | | |
| Tu 2/12 | Game Trees: Expectimax; Utilities | Ch. 5.2-5 (2e: Ch. 6.2-5) | 2PP · 6PP | video | HW3 | 2/18 |
| | | | | | P2: Multi-Agent Pacman | 2/22 5pm |
| Th 2/14 | Markov Decision Processes | Ch. 16.1-3 | 2PP · 6PP | video | | |
| Tu 2/19 | Markov Decision Processes II | Sutton and Barto Ch. 3-4 | 2PP · 6PP | video | HW4 | 2/25 |
| Th 2/21 | Reinforcement Learning | Ch. 17.1-3, S&B Ch. 6.1,2,5 | 2PP · 6PP | video | | |
| Tu 2/26 | Reinforcement | | 2PP · 6PP | video | HW5 | 3/4 |

| Day | Topic | Reading | Slides | Video | Assignment | Due |
|---|---|---|---|---|---|---|
| | Learning II | | | | | |
| | | | | | P3: Reinforcement Learning | 3/8 5pm |
| Th 2/28 | Probability | Ch. 13.1-5 (2e: Ch. 13.1-6) | 2PP · 6PP | video | | |
| Tu 3/5 | Bayes' Nets: Representation | Ch. 14.1-2,4 | 2PP · 6PP | video | Practice Midterm | 3/10 |
| Th 3/7 | Bayes' Nets: Independence | Ch. 14.1-2,4 | 2PP · 6PP | video | | |
| Tu 3/12 | MIDTERM 1 (time/place TBA) (no lecture) | | | | HW6 | 3/18 |
| Th 3/14 | Bayes' Nets: Inference | Ch. 14.3, Jordan 2.1 | 2PP · 6PP | video | | |
| Tu 3/19 | Bayes' Nets: | Ch. 14.4-5 | 2PP · 6PP | video | HW7 | 4/2 |

| Day | Topic | Reading | Slides | Video | Assignment | Due |
|---|---|---|---|---|---|---|
| | Sampling | | | | | |
| Th 3/21 | Decision Diagrams / VPI | Ch. 15.1-3, 6 | 2PP · 6PP | video | | |
| Tu 3/26 | Spring Break | | | | | |
| Th 3/28 | Spring Break | | | | | |
| Tu 4/2 | HMMs: Filtering | Ch. 15.2,5 | 2PP · 6PP | video | HW8 | 4/8 |
| | | | | | P4: Ghostbusters | 4/12 5pm |
| Th 4/4 | HMMs: Wrap-up / Speech | Ch. 15.2,6 | 2PP · 6PP | video | | |
| Tu 4/9 | ML: Naive Bayes | Ch. 15.2,6 | 2PP · 6PP | video | HW9 | 4/15 |
| Th 4/11 | ML: Perceptrons | Ch. 15.2,6 | 2PP · 6PP | video | Contest: Pacman Capture the Flag | 5/1 |

| Day | Topic | Reading | Slides | Video | Assignment | Due |
|---|---|---|---|---|---|---|
| Tu 4/16 | ML: Kernels and Clustering | | 2PP · 6PP | video | Practice Midterm 2 | 4/19 |
| Th 4/18 | ML: Decision Trees and Neural Networks | | 2PP · 6PP | video | | |
| Tu 4/23 | MIDTERM 2 (time/place TBA) (no lecture) | | | | P5: Classification | 5/3 5pm |
| Th 4/25 | Robotics / Language / Vision | | 2PP · 6PP | video | | |
| Tu 4/30 | Robotics / Language / Vision | | 2PP · 6PP | video | | |
| Th 5/2 | Advanced Topics and Final Contest | | 2PP · 6PP | video | Practice Final | 5/8 |
| F 5/17 | FINAL EXAM (11:30-2:30pm) | | | | | |

# Course Information

## TIMES AND LOCATIONS

- **Lecture:** Tuesdays and Thursdays 5-6:30pm in 2050 VLSB
- **Sections:** Beginning the week of 1/28
  - 101: Th 9-10am, 85 Evans (Jeff)
  - 102: Th 10-11am, 75 Evans (Jeff)
  - 103: Th 11-noon, 285 Cory (James)
  - 104: Th 1-2pm, 85 Evans (Judy)
  - 105: Th 2-3pm, 71 Evans (Judy)
  - 106: Th 3-4pm, 101 Wheeler (Brad)
  - 107: F 10-11am, 101 Wheeler (Nick)
  - 108: F 11-noon, 101 Wheeler (James)

## COMMUNICATION

There will be several routes of communication for this course:

- Announcements will be posted to this website.

- The main mode of electronic communication between students and staff, as well as amongst students, will be through Piazza. It is intended for general questions about the course, clarifications about assignments, student questions to each other, discussions about material, and so on. We strongly encourage students to participate in discussion, ask, and answer questions through this site. The course staff will monitor discussions closely.

- If you need to contact the course staff privately, you should email cs188-staff AT lists.berkeley.edu. You may of course contact the professor or GSIs directly, but the staff list will produce the fastest response.

## COURSE DESCRIPTION

This course will introduce the basic ideas and techniques underlying the design of intelligent computer systems. A specific emphasis will be on the statistical and decision-theoretic modeling paradigm.

By the end of this course, you will have built autonomous agents that efficiently make decisions in fully informed, partially observable and adversarial settings. Your agents will draw inferences in uncertain environments and optimize actions for arbitrary reward structures. Your machine learning algorithms will classify handwritten digits and photographs. The techniques you learn in this course apply to a wide variety of artificial intelligence problems and will serve as the foundation for further study in any application area you choose to pursue.

See the syllabus for slides, deadlines, and the lecture schedule.

## PREREQUISITES

- **CS 61A or 61B:** Prior computer programming experience is expected (see below); most students will have taken both these courses.
- **CS 70 or Math 55:** Facility with basic concepts of propositional logic and probability are expected (see below); CS 70 is the better choice for this course.

This course has substantial elements of both programming and mathematics, because these elements are central to modern AI. You should be prepared to review basic probability on your own if it is not fresh in your head. You should also be very comfortable programming on the level of CS 61B even though it is not strictly required.

**CS61A AND CS61B AND CS70 is the recommended background.**

## LANGUAGE

Course programming assignments will be in Python. We do not assume that students have previous experience with the language, but we do expect you to learn the basics very rapidly. Project 0 is designed to teach you the basics of Python, but if you want to get a head start here is a good tutorial: ACM Python Tutorial

## ASSIGNMENTS

This class includes **five programming projects** and **regular electronic assignments**.

**Slip days:** Programming projects must be turned in electronically by 5:00pm on the listed due date. You will have **5 slip days** for these projects, up to **two of which can be used for each project**. Note that slip days are counted at the granularity of days, rounded up to the nearest day. The other homework assignments do not have late days!

**Ethics:** Submissions should acknowledge all collaborators and sources consulted. All code and written responses should be original. we trust you all to submit your own work, but to protect the integrity of the course from anyone who doesn't want to play by the rules, we will actively be checking for code plagiarism (both from current classmates and previous semesters). I'm not lenient about cheating; I sympathize with Kris Pister's policy.

## GRADING

Overall grades will be determined from:

- Programming Assignments (30%)
- Electronic Assignments (10%)
- Midterm exam 1 (10%)
- Midterm Exam 2 (15%)
- Final Exam (35%)

Grades are on the following fixed scale

A [90 – 100]%
A- [85 – 90)%
B+ [80 – 85)%
B [75 – 80)%
B- [70 – 75)%
C+ [65 – 70)%
C [60 – 65)%
C- [55 – 60)%
D+ [50 – 55)%
D [45 – 50)%
D- [40 – 45)%
F [0 – 40)%

These cutoffs represent grade minimums. We may adjust grades upward based on class participation, extra credit, etc. The grade of A+ will be awarded at the professor's discretion based on exceptional performance.

**Regrade Policy:** If you believe an error has been made in the grading of one of your exams or assignments, you may resubmit it for a regrade. For any regrade requests, send an email to the staff list with a detailed explanation of which problems you think we marked incorrectly and why. Regrades for cases where we mis-applied a rubric in an individual case are much more likely to be successful than regrades that argue about relative point values within the rubric, as the rubric is applied to the entire class. Because we will examine your entire submission in detail, your grade can go up or down as a result of a regrade request.

## TEXTBOOK

No textbook is required for this course. For students who want to read more, we recommend Artificial Intelligence: A Modern Approach by Stuart Russell (UC Berkeley) and Peter Norvig (Google). Be aware, however, that this is not a course textbook, so our presentation does not necessarily follow the presentation in the book.

## ENROLLMENT

Here are the policies that govern admission into classes, and here are some answers to frequently asked questions about admission. The course staff does not control enrollment!