



⚠ To see course content, [sign in](#) or [register](#).

Course Policies

Times and Locations

- **Lecture:** Mondays and Wednesdays 5:00-6:30pm in 2050 Valley Life Sciences
- **Discussion Sections:** come back soon for the schedule!
- **Exam Practice Sessions:** come back soon for the schedule!

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](mailto:cs188-staff@lists.berkeley.edu). You may of course contact the professor or GSIs directly, but the staff list will produce the fastest response.
- For exceptions (projects, homework, exams) please email Davis Foote, the

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 and 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 **six programming projects** and **regular electronic assignments**.

Collaboration: Project 0 is to be completed alone. Projects 1 through 6 can be completed alone or in teams of two. When completing in a team of two it is important that both team members submit their team's project files into edX to ensure a grade is registered for both team members. Homework is to be submitted individually but may be discussed in groups.

Slip days: Programming projects must be turned in electronically by 5:00pm on the listed due date. You will have a total of **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. For example, for a project due at 5pm on Friday, any submission from Friday 5pm - Saturday 5pm will use up one slip day, any submission from Saturday 5pm - Sunday 5pm will use up two slip days, and any submissions after Sunday 5pm will receive no credit. **The other homework assignments do not have late days!**

Project grading: Projects will by default be graded automatically for correctness, though we will review projects individually as necessary to ensure that they receive the credit they deserve. To calculate your total project grade, we will consider all possible allocations of slip days across all projects to give you the best score achieved. For example, suppose you submitted a 24/25 on Project 1 on time, and then resubmitted a 25/25 on Project 1 after using two slip days. If by Project 5, you used up all your slip days, but you need to use an additional slip day on Project 5 and wish you had not submitted two days late for the one

point in Project 1, you can submit into the late module for Project 5 (edX does not block you from going over the number of slip days) and we will automatically redistribute your slip days usage.

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](#).

Exams

There will be no make-up exams.

Grading

Overall grades will be determined from:

- Programming Assignments (30%)
- Electronic Assignments (10%)
- Midterm Exam (25%)
- 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. 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!