

Syllabus

UPDATED 24 DAYS AGO

If you have any questions after reading this Syllabus, post on our Ed discussion forum!

Teaching Team



Kelly Cochran
kcochran @ stanford dot edu
Co-Instructor



Joel Ramirez
joel101 @ stanford dot edu
Co-Instructor



Isabel Michel
imichel @ stanford dot edu
TA

I. Course Overview

While the initial foundations of computer science began in the world of discrete mathematics (after all, modern computers are digital in nature), recent years have seen a surge in the use of probability as a tool for the analysis and development of new algorithms and systems. As a result, it is becoming increasingly important for budding computer scientists to understand probability theory, both to provide new perspectives on existing ideas and to help further advance the field in new ways.

CS109: Probability for Computer Scientists starts by providing a fundamental grounding in combinatorics, and then quickly moves into the basics of probability theory. We will then cover many essential concepts in probability theory, including particular probability distributions, properties of probabilities, and mathematical tools for analyzing probabilities. Finally, the last third of the class will focus on data analysis and machine learning as a means for seeing direct applications of probability in this exciting and quickly growing subfield of computer science. This is going to be a great quarter and we are looking forward to the chance to teach you.

Learning Goals

Our goal in CS109 is to build foundational skills and give you experience in the following areas:

1. Understanding the combinatorial nature of problems: Many real problems are based on understanding the multitude of possible outcomes that may occur, and determining which of those outcomes satisfy some criteria we care about. Such understanding is important both for determining how likely an outcome is, but also for understanding what factors may affect the outcome (and which of those may be in our control).
2. Working knowledge of probability theory: Having a solid knowledge of probability theory is essential for computer scientists today. Such knowledge includes theoretical fundamentals as well as an appreciation for how that theory can be successfully applied in practice. We hope to impart both these concepts in this class.
3. Appreciation for probabilistic statements: In the world around us, probabilistic statements are often made, but are easily misunderstood. For example, when a candidate in an election is said to have a 53% likelihood of winning does this mean that the candidate is likely to get 53% of the vote, or that that if 100 elections were held today, the candidate would win 53% of them? Understanding the difference between these statements requires an understanding of the model in the underlying probabilistic analysis.
4. Applications: We are not studying probability theory simply for the joy of drawing summation

symbols (okay, maybe some people are, but there are even more exciting things you can do with probability than that!), but rather because there are a wide variety of applications where probability allows us to solve problems that might otherwise be out of reach (or would be solved more poorly without the tools that probability can bring to bear). We'll look at examples of such applications throughout the class.

5. An introduction to machine learning: Machine learning is a quickly evolving subfield of artificial intelligence which has grown to impact many applications in computing. It focuses on analyzing large quantities of data to build models that can then be harnessed in real problems, such as filtering email, improving web search, understanding computer system performance, predicting financial markets, or analyzing DNA.

Course Topics

Here are the broad strokes of the course. More information is available on our [Schedule](#) page. We cover a very broad set of topics so that you are equipped with the probability and statistics you will see in your future CS studies!

- Counting and probability fundamentals
- Single-dimensional random variables
- Probabilistic models
- Uncertainty theory
- Parameter estimation
- Introduction to machine learning

Prerequisites

The prerequisites for this course are CS103, CS106B or X, and Math 51 (or equivalent courses). Probability involves a fair bit of mathematics (set theory, calculus, and familiarity with linear algebra), and we'll be considering several applications of probability in CS that require familiarity with algorithms and data structures covered in CS106B/X. Here is a quick rundown of some of the mathematical tools from CS103 and Math 51 that we'll be using in this class: multivariate calculus (integration and differentiation), linear algebra (basic operations on vectors and matrices), an understanding of the basics of set theory (subsets, complements, unions, intersections, cardinality, etc.), and familiarity with basic proof techniques. We'll also do combinatorics in the class, but we'll be covering a fair bit of that material ourselves in the first week. Past students have managed to take CS106B concurrently with CS109 and have done just fine. CS103 is the pre-requisite that we rely on the least. Students have done well even without having taken CS103.

II. Course Structure

Lectures

Lectures are MWF from 3:00pm until 4:15pm. We will be holding live lectures in-person in NVIDIA Auditorium, in the basement of Huang in the Engineering Quad. Come to learn the material and engage in interesting problems collectively with the class. While lecture attendance isn't mandatory, we will give extra credit for attendance, and it is correlated with doing well in the course and mastering the material. Lecture is a good place to learn, make friends and enjoy probability. Good times!

Lecture Recordings

Lectures will be recorded, as we are an SCPD class this quarter. You can access these recordings through Canvas.

Units

If you are a Stanford undergraduate, you are **required** to take CS109 for 5 units of credit (this is by department and university policy, no exceptions). If you are a graduate student, you may enroll in CS109 for 3 or 4 units if it is necessary for you to reduce your units for administrative reasons. Taking the course for reduced units does not imply any change in the course requirements.

Sections

Active participation plays an important role in making you adept at combining probability and computer science. It has also been observed over many quarters that keeping up with the material highly correlates with improved class performance.

Each week, for one hour, you will meet in a small group with a TA and work through problems. If you have taken any of the CS106 classes, our sections will be very similar—except with more probability. The sign-up form for sections will go out on the first day of class, and the deadline for filling out that form is midnight on Wednesday night. We will let you know which section you are in by the weekend, and you will have your first section during Week 2. You do not need to enroll in your section via Axxess.

Section attendance and participation is required for all students. Your participation grade is based on how much you engage with section; a student that engages and helps others will receive a better score than one who shows up late or doesn't participate. Students are allowed **two makeup sections** (i.e. attending another TA's section in the same week) and **two section absences** in the quarter without penalty.

Grading

The grade for the course will be determined according to the following breakdown:

Component	Final grade
Problem Sets	40%
Midterm	20%
Final	30%
Section Participation	10%
Lecture Attendance (Extra Credit)	3%

Problem Sets

During the course, there will be five problem sets assigned. We put a lot of love into these problems so that they can help train you to become gifted practitioners of probability and computer science. Use them as practice. Doing well on the problem sets is the best way to prepare for life after CS109 (and the exams). Each student is to submit individual work on the problem sets. The problem sets will often include coding tasks using Python. We also strongly encourage you to learn LaTeX, which is the interchangeable markup language for typing math on a computer. We will hold review sessions for those of you who are not familiar with python.

We use "the psetapp", a webapp written by Chris Piech, for all homework in CS109. Importantly, the psetapp allows you to check your work as you go! That way, you can get immediate feedback as to whether you have figured out the solution. Things to know: 1) The psetapp auto-saves your progress, so you never need to hit a "save" or "submit" button. At the deadline, you will no longer be able to edit your work, and we will start grading your current saved solutions. (2) You can check your answer as many times as you want. There is no penalty for retrying if you didn't get the answer correct the first time.

After you submit, we will use a combination of manual and automatic grading. For every problem, we expect that you show your work, meaning that getting full points for each problem requires you to explain how you got your answer. The process of articulating how you arrived at a correct solution can be just as helpful for learning as getting to the solution in the first place, so we encourage you to take this component of the problem sets seriously. Never write down just the answer. For each math solution, you should provide a detailed enough explanation that someone who is fluent in CS109 would understand how to solve it. This could look like a written paragraph of text, or a series of math steps written out with a bullet point of explanation for each step. For coding solutions, we expect you to use enough comments and good style that it is clear what process your code is using to solve the task. Coding solutions are not graded on "efficiency"; that is, you do not need to minimize the number of lines of code you use or how long your code takes to run, within reason (i.e. no CS109 problem should need 1000+ lines of code to solve).

After grades are released, you have **one week** to file a regrade request if you think that points were deducted incorrectly by emailing the teaching team. We reserve the right to regrade your entire pset.

Late Policy

There may be unforeseen circumstances that make it difficult to turn in homework assignments on time. Our philosophy is to treat you as adults and thus we have a generous late policy to reflect the many different needs that may come up. However, the course will end for everyone on the same date. As such, if you are late on one problem set, you will have to work extra hard to catch up. Time management can be hard and we encourage you to give it the full respect it deserves. In practice, falling behind often impacts midterm and final exam scores.

- **Due Date:** The on-time deadline is listed on each pset's homepage in the psetapp (generally at 2 pm Pacific Time). This deadline is computer-enforced, so there is no wiggle room. Finishing the assignment by the deadline means that you are in sync with the course. Hooray!
- **Grace Period:** For each pset deadline, you have the option of taking a "grace period" mini-extension. This grace period lasts an hour and is meant to give you the flexibility needed to account for issues like spotty wifi or forgetting to write out one of your explanations before the original deadline. You can take as many grace period extensions as needed, with no impact to your grade, via the psetapp extension request button.
- **Extension:** CS109 is a fast-paced class, and we encourage you to aim for meeting the deadline for each pset to avoid falling behind on future problem sets (or the midterm / final). Having said that, you might have a medical, personal or serious time-management situation which requires an "as long as possible" extension. We will give you up until the date when we release solutions. This time is set on a per-pset basis but generally is two class days after the problem set was due. For example, if the pset was due on Friday, typically an extension would allow you to make changes until the following Wednesday. This is the longest extension we can offer. You can grant yourself an extension via the psetapp, so you do not need to receive permission from the teaching team, and you do not need to worry about providing a "serious enough" reason for taking the extension. However, to encourage students to avoid falling behind, the maximum number of extensions you can take without it impacting your grade is two. Once you have already used two extensions, the highest grade you can obtain on psets where you take further extensions past two will be capped at 85%.
- **After the Hard Extension Deadline:** What if you are not done by the time we release solutions? In general we do not accept work after the solutions have been posted and TAs start grading. As such, you should make sure you don't accidentally miss this very hard deadline! Having said that, there may be a real crisis that means you are not able to do your work before the solutions are released (e.g. a hospital stay, funeral attendance, etc). First, we hope you are well. Personal life is so truly important and we respect you doing what you need to do. In such an extreme case, you need to contact the course staff (cs109@cs.stanford.edu) as well as your [Undergraduate Advising Director](#), if you are a Stanford student, and we will work something out. Please do contact us as early as possible.

Warning: Pset3 (the pset before the midterm) will have a shorter extension window than other psets. Why? The pset is due right before the midterm, and we need to release the solutions before the midterm so that students can use them while studying. Similarly, Pset5 (the pset before the final) will have an extension deadline which reflects when we need to release solutions.

Exams

In addition to the assignments, there will be a midterm and a final exam:

- Midterm: Tuesday, July 23rd, 7 - 9pm
- Final: Sat, Aug 17th, 3:30 - 6:30pm

Midterm Alternatives: A week before the midterm, we will give more information on what you should do if you have an academic conflict with the midterm.

No Final Alternatives: For a variety of reasons (including university policy), we generally do not provide alternate offerings for the final exam. Please make sure that you can attend the final exam at the specified time before enrolling in the class. Typically, you will not be able to make the final if you are taking another class at the same time as CS109. If you cannot make the final exam, you should plan on taking CS109 in a different quarter. If you have a particularly compelling reason for missing the exam, and you contact us the first week of the term, we will look into making an exception to this generally hard set rule. If you start the course and for some reason you can't take the exam, we will offer you an incomplete and you can take the final with a future offering of CS109.

Lecture Attendance

Attending lecture is correlated with students keeping up with the pace of the course and more successful outcomes in CS109. So, to encourage you to attend lectures, we will offer a small amount of extra credit (3%) for consistent attendance, which will be incorporated into your final grade. To get credit for attendance each lecture, you will need to answer the short concept check question in the psetapp. For non-SCPD students, the deadline to get full credit for each concept check is an hour after lecture (so 5:15pm), and we will give you time in lecture to submit your answer; we will still award partial credit for submitting within 48 hours. For SCPD (remote) students, the deadline is 48 hours after lecture to account for lecture upload speed and timezone differences.

Extra Credit Challenge

Each quarter CS109 holds an extra credit challenge, where students apply the principles in this class to explore a topic of their own choice. Participation in the contest is completely optional, and prizes involve extra credit to your final course grade. More details about the contest to be posted midway through the quarter.

III. Course Resources

CS109 Course Reader

We have a beautiful CS109 Course Reader, created by Chris Piech, that you can use as reference throughout the course. It is like a textbook, but more condensed; not every detail covered in lecture is in the course reader, but it contains helpful summaries of all the main concepts and some great worked-through example problems.

Optional Textbook

Sheldon Ross, *A First Course in Probability* (10th Ed.), Pearson Prentice Hall, 2018.

This is an *optional* textbook, meaning that the text is not required material, but students may find Ross offers a different and useful perspective on the important concepts of the class. The 8th, 9th, and 10th editions of the textbook are all fine for this class.

Borrowing the textbook online: HathiTrust, a library archive of which Stanford is a member, has granted the university online access to the 8th edition (2010) for the duration of the quarter. The "check out" system works similarly to print reserves: A user can check out the book an hour at a time as long as they are actively using it. Access guidelines are on the [HathiTrust How To Use It webpage](#). Once you're logged in, the book is at [this link](#).

"Working" Office Hours

To help you master the material in this class, the course staff will hold "working" office hours. The idea is that you can come work on your problem sets at these office hours, so you can immediately ask any questions that come up while working on your problem sets. You don't need to have a question in mind already to come to office hours at all! While you are certainly not required to attend any of these working office hours, they are simply meant to encourage you to interact with the course staff more often in order to help you better understand the course material. Besides, our job is to help everyone learn the material for this class, and being more accessible to you when you are actually working on your assignments (rather than when you just have a problem) will help the course go more smoothly for you (and it'll be more fun for us).

Not all office hours are the same! [The office hour calendar](#) will indicate specific details such as the time and location of each office hours. Office hours are either in person or online. If an office hour time on the calendar says it is in person, there will be no online option — and if the calendar says an office hours session is online, please don't show up at our TAs' dorms :).

Accommodations

Students who may need an academic accommodation based on the impact of a disability must initiate the request with the Office of Accessible Education (OAE). Professional staff will evaluate the request with required documentation, recommend reasonable accommodations, and prepare an Accommodation Letter for faculty. For students who have disabilities that don't typically change appreciably over time, the letter from the OAE will be for the entire academic year; other letters will be for the current quarter only. Students should contact the OAE as soon as possible since timely notice (for example, at least a week before an exam) is needed to coordinate accommodations. Students should also send your accommodation letter to instructors as soon as possible.

IV. Honor Code

Please read our full [Honor Code Policy](#), which specifically prohibits you from soliciting or taking solutions from other students, websites like Stack Overflow or Chegg, and ChatGPT.

Looking Forward to a Great Quarter

Genuinely, teaching CS109 is a profound joy. Thanks for coming to learn with us. We can't wait 🌱.