

CS109 Course Overview

UPDATED A MONTH AGO

If you have any questions after reading this overview, head on over to our [FAQ page](#), post on our [discussion forum](#), or email us at cs109@cs.stanford.edu.

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 common probability distributions, properties of probabilities, and mathematical tools for analyzing probabilities. The last third of the class will focus on data analysis and machine learning as direct applications of probability we've learned in the weeks prior. **Read more [here](#) to learn what CS109 is all about.** This is going to be a great quarter and we are looking forward to the chance to teach you!!

Teaching Team







Lecturer: Jerry Cain

✉ jerry@cs.stanford.edu

🕒 Office Hours: Wednesdays, 4:30 - 7:00pm, Huang Basement

Superstar CS109 CAs!



I. Course Overview

Prerequisites

The prerequisites for this course are CS103, CS106B, 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. 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. See our [FAQ](#) for more information.

Learning Goals

After you're finished with CS109, we hope you'll have achieved the following learning goals:

- Reason about situations using probabilities, expectation, and variance.
- Feel comfortable learning about new probability concepts beyond the scope of this class (e.g., through one's own research, studies, or interests).
- Write programs to simulate random experiments and to test experiment hypotheses.
- Understand and implement simple machine learning algorithms like Naive Bayes and Logistic Regression.

Course Topics

Here are the broad strokes of the course (in approximate order). 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

II. Course Structure

Units

If you are an 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.

Lectures

We will be holding live lectures MWF at 3:00pm, and I'll lecture for no more than 80 minutes, and often less. (The 4:20pm end time is a limit, but not a goal.) Come to learn the material and engage in interesting problems collectively with the class. While lecture attendance isn't mandatory, it is correlated with doing well in the course and mastering the material. If you can't make lecture time, you can watch the lectures on Canvas.

Sections

Active participation plays an important role in making you a master of probability for 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 50 minutes you will meet in a small group with one of our outstanding CAs and work through problems. If you have taken any of the CS106 classes, our sections will be very similar—except with more probability.

Grading

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

Component	Final grade
Concept Checks	10%
Weekly Problem Sets (6 x 8%)	48%
Mid-Quarter Exams (2 x 11%)	22%

Component	Final grade
Final Exam	15 - 20%
Section Participation	0 - 5%

Concept Checks

In CS109 you will learn a large amount of content. The class moves at a consistent pace and staying on top of the material is key to success in the class and mastering the material. To help you keep pace, we will post a concept check by 9:00pm on Monday, Wednesday, and Friday evenings. Concept checks require you answer a small number of questions exercising the material introduced during the corresponding lecture.

Recognizing that many of you will prefer to watch or rewatch lecture videos over the course of a weekend, we'll allow all three of a week's concept checks to be submitted just before lecture starts the following Monday. Restated, Week 1's concept checks will fall due at 3:00pm on the Monday of Week 2, Week 2's concept checks will fall due at 3:00pm on the Monday of Week 3, and so forth. Of course, you benefit by submitting the concept checks as quickly as possible, but we also understand that sometimes you just need to ignore a class for a few days and catch up later on.

Unless you're enduring a family emergency or a severe illness, you can't submit concept checks after the Monday deadline. In particular, we can't extend deadlines just because you're busy or you forgot. These concept checks don't serve their purpose of keeping you on schedule if we allow the deadlines to slip.

Problem Sets

During the course, there will be six problem sets assigned. **Each student is to submit individual work for the problem sets**. You may discuss with other students and course staff, but you must cite all discussion on your individually written final write-up of the problem set. All homework assignments should be turned in at the beginning of class on the published due date.

Late Policy (Problem Sets)

We anticipate that there may be unforeseen circumstances that make it difficult to turn in homework assignments on time. As opposed to the concept checks, the problem sets require a serious time investment and sometimes they just don't come together by the published deadline.

- **All students** will be granted a penalty-free "**grace period**" for submission on all problem sets (except the final assignment). The grace period is one full **class day** and allows you to submit the assignment after the original deadline, with no impact on the final grade. As an example, a problem set due on Monday at 3:00pm may be turned in on Wednesday before 3:00pm to submit for full credit, or problem set due on Friday at 3:00pm may be turned in on the following Monday before 3:00pm Pacific to submit for full credit. This grace period is meant to give built-in flexibility for any unexpected snags—however, we strongly recommend that students submit by the original deadline if possible, in order to avoid falling behind.
- You are permitted to grant yourself *additional* penalty-free extensions—beyond even the first self-granted extension—as needed, though you need to email Jerry and Head CA Will Song to let them know you need the additional time. In general, you are never penalized for lateness provided you submit an assignment within a week of the original due date.

Sections

Section attendance and participation is required of all students. Additionally, students are allowed two "makeup sessions" in the quarter without penalty. A makeup session is defined as one which is not their originally assigned section. Students are allowed one unexcused absence in the quarter without penalty. Understand that absences can only be excused because of severe illness, COVID isolation, family emergency, or Stanford-sponsored travel.

Your section grade is automatically a 100%. If you attend all sections, then you're awesome and your section grade counts 5% toward your final grade. For every section you miss, your section grade counts 1% less and your final exam grade counts 1% more. If you miss five or more sections, then your section grade doesn't contribute to your final grade.

Exams

In addition to the assignments, there will be two in-class exams during Weeks 4 and 7, and an in-class final exam during our regularly scheduled final exam slot.

- Week 4 Exam: Thursday, February 2nd, 7:00 - 9:00pm
- Week 7 Exam: Thursday, February 23rd, 7:00 - 9:00pm
- Final Exam: Tuesday, March 21st, 8:30 - 11:30am

Each of the two mid-quarter exams are traditional, pencil-on-paper exams, save for the fact that they are held at night so that we can give you more time to complete them. Each of the two mid-quarter exams is closed-book, closed-calculator, closed-computer, although you are permitted to prepare and bring **two** pages—front and back—of notes with you and refer to them during the exam.

Exam material will cover the material you have learned in the class lectures, class handouts, concept checks, and then assignments.

The final exam will be a traditional in-class, pencil on paper, closed-book, closed-calculator, closed-computer exam, though you're permitted to bring **four** double-sided pages of notes instead of just two. The final exam is scheduled for Tuesday, March 21st from 8:30 - 11:30am. For those with a legitimate conflict—e.g., you're taking another class at the same time and it's offering a final exam or requiring a final presentation during the same time slot, an alternate final exam will be offered the same day from 12:15pm - 3:15pm.

Extra Credit Contest

CS109 has traditionally held an extra credit contest, 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

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. Suggested, optional reading assignments from the textbook (10th Ed.) are in the schedule on the course website. 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 Fall 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](#).

All students should retain receipts for books and other course-related expenses, as these may be qualified educational expenses for tax purposes. If you are an undergraduate receiving financial aid, you may be eligible for additional financial aid for required books and course materials if these expenses exceed the aid amount in your award letter. For more information, review your award letter or visit the [Student Budget website](#).

"Working" Office Hours

To help make you more successful in this class, the course staff will hold "working" office hours. The idea is to encourage you to 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. 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).

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. **If you require additional, or different, accommodations specific to the Spring 2022 learning environment, please contact your OAE adviser.**

IV. Honor Code

Each student is expected to submit their own work on the CS109 problem sets. Students may discuss problem sets with each other as well as the course staff, and you can even confirm a final numeric result so you know that you've computed the correct answer. You may not, however, present another person's reasoning behind a proof or result as your own. **Any discussion of problem set questions with others should be noted on a student's final write-up of the problem set answers.** Each student must turn in their own solution. Excessive collaboration can result in honor code violations. Questions regarding acceptable collaboration should be directed to the class instructor prior to the collaboration.

It is a violation of the honor code to copy problem set from others, or to copy or derive them from solutions found online or in textbooks, previous instances of this course, or other courses covering the same topics (e.g., Stats 116 or probability courses at other schools). Copying of solutions from students who previously took this or a similar course is also a violation of the honor code. Finally, a good point to keep in mind is that you must be able to explain and/or re-derive anything that you submit.

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

Looking Forward to a Great Quarter

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