CS103: Mathematical Foundations of Computer Science, Stanford University

and

CS109: Probability Theory for Computer Scientists, Stanford University

Stanford, CA, USA
Keith Schwarz and Mehran Sahami
{htiek, sahami}@cs.stanford.edu
Course URLs:
cs103.stanford.edu
cs109.stanford.edu

Knowledge Areas that contain topics and learning outcomes covered in the course

Knowledge Area	Total Hours of Coverage
Discrete Structures (DS)	30
Algorithms and Complexity (AL)	6
Intelligent Systems (IS)	2

Where does the course fit in your curriculum?

CS103 and CS109 make up the first two courses in the required introductory CS theory core at Stanford. The prerequisites for CS103 are CS2 and mathematical maturity (e.g., comfortable with algebra, but calculus is not a requirement). The prerequisites for CS109 are CS2, CS103, and calculus. However, calculus is only used for topics beyond the CS2013 Discrete Structures guidelines, such as working with continuous probability density functions. Approximately 400 students take each course each year. The majority of students taking the courses are sophomores, although students at all levels (from freshman to graduate students) enroll in these courses.

What is covered in the course?

CS103 covers:

- Sets
- Functions and Relations
- Proof techniques (including direct, contradiction, diagonalization and induction)
- Graphs
- Logic (proposition and predicate)
- Finite Automata (DFAs, NFAs, PDAs)
- Regular and Context-Free Languages
- Turing Machines
- Complexity Classes (P, NP, Exp)
- NP-Completeness

CS109 covers:

- Counting
- Combinations and Permutations
- Probability (including conditional probability, independence, and conditional independence)
- Expectation and Variance
- Covariance and Correlation
- Discrete distributions (including Binomial, Negative Binomial, Poisson, and Hypergeometric)
- Continuous distributions (including Uniform, Normal, Exponential, and Beta)
- Limit/Concentration results (including Central Limit Theorem, Markov/Chebyshev bounds)
- Parameter estimation (including maximum likelihood and Bayesian estimation)
- Classification (including Naive Bayes Classifier and Logistic Regression)
- Simulation

What is the format of the course?

Both CS103 and CS109 use a lecture format, but also include interactive class demonstrations. Each course meets three times per week for 75 minutes per class meeting. CS103 also offers an optional 75 minute discussion session. The courses each run for 10 weeks (Stanford is on the quarter system).

How are students assessed?

CS103 currently requires nine problem sets (approximately one every week), with an expectation that students spend roughly 10 hours per week on the assignments. The problem sets are comprised of rigorous exercises (e.g., proofs, constructions, etc.) that cover the material from class during the just completed week.

CS109 currently requires five problem sets and one programming assignment (one assignment due every 1.5 weeks), with an expectation that students spend roughly 10 hours per week on the assignments. The problem sets present problems in probability (both applied and theoretical) with a bent toward applications in computer science. The programming assignment requires students to implement various probabilistic classification techniques, apply them to real data, and analyze the results.

Course textbooks and materials

CS103 uses two texts (in addition to a number of instructor-written course notes):

1. Chapter One of <u>Discrete Mathematics and Its Applications</u>, by Kenneth Rosen. This chapter (not the whole text) covers mathematical logic.
2. <u>Introduction to the Theory of Computation</u> by Michael Sipser.

CS109 uses the text <u>A First Course in Probability Theory</u> by Sheldon Ross for the first two-thirds of the course. The last third of the course relies on an instructor-written set of notes/slides that cover parameter estimation and provide an introduction to machine learning. Those slides are available here:
http://ai.stanford.edu/users/sahami/cs109/

Why do you teach the course this way?

As the result of a department-wide curriculum revision, we created this two course sequence to capture the foundations we expected students to have in discrete math and probability with more advanced topics, such as automata, complexity, and machine learning. This obviated the need for later full course requirements in automata/complexity and an introduction to AI (from which search-based SAT solving and machine learning were thought to be the most critical aspects). Students do generally find these courses to be challenging.

Body of Knowledge coverage

KA	Knowledge Unit	Topics Covered	Hours
DS	Proof Techniques	All	7
DS	Basic Logic	All	6
DS	Discrete Probability	All	6
AL	Basic Automata, Computability and Complexity	All	6
DS	Basics of Counting	All	5
DS	Sets, Relations, Functions	All	4
DS	Graphs and Trees	All Core-Tier1	2
IS	Basic Machine Learning	All	2

Additional topics

CS103 covers some elective material from:

AL/Advanced Computational Complexity
AL/Advanced Automata Theory and Computability

CS109 provides expanded coverage of probability, including:

Continuous distributions (including Uniform, Normal, Exponential, and Beta)
Covariance and Correlation
Limit/Concentration results (including Central Limit Theorem, Markov/Chebyshev bounds)
Parameter estimation (including maximum likelihood and Bayesian estimation)
Simulation of probability distributions by computer

CS109 also includes some elective material from:

IS/Reasoning Under Uncertainty
IS/Advanced Machine Learning

Other comments

Both these courses lectures move quite rapidly. As a result, we often cover the full set of topics in one of the CS2013 Knowledge Units in less time than proscribed in the guidelines. The "Hours" column in the Body of Knowledge coverage table reflects the number of hours we spend in lecture covering those topics, not the number suggested in CS2013 (which is always greater than or equal to the number we report).