Syllabus

1 Logistics

- Instructor: Tselil Schramm (tselil@stanford.edu). Office Hours TBD.
- Course website: www.tselilschramm.org/stochproc/stats217-winter23.html
- Teaching Assistants: Isaac Gibbs and Nathan Tung
- Lectures: Tuesdays & Thursdays, 10:30-11:50, Huang 18.

The best way to contact course staff is by email. Please be sure to include "STATS217" in the subject line.

Prerequisites: Students will need to have taken probability theory (STATS116, CS109, or MATH151), which in turn necessitates a background in linear algebra (MATH51 or CS205) and calculus (MATH19, MATH20, MATH21). If you have not taken the courses listed but believe that you have sufficient background to take the course, please contact me.

2 Overview

This course is an introduction to discrete stochastic processes. We will see how to model real-world stochastic processes as simple, structured random systems, and how doing so gives us the power to draw remarkably precise, controlled conclusions about the macroscopic behavior of these chaotic processes. Topics covered include discrete and continuous time Markov Chains, Martingales, Poisson Processes, and some topics in Statistical physics. This is a rigorous, proof-based course, but we will not require knowledge of measure theory.

Course Description (as it appears in the course catalogue): Discrete and continuous time Markov chains, Poisson processes, random walks, branching processes, first passage times, recurrence and transience, stationary distributions. Non-Statistics masters students may want to consider taking STATS 215 instead.

Topics We will cover the following topics, roughly partitioned into units:

1. **Discrete-time Markov chains**. A Markov chain is a discrete-time random process in which each step is independent of the past, conditioned on the present. This class of processes is very expressive and can be used model random walks, gambling and games, as well give basic models for weather and the spread of infectious diseases. Besides for modeling, algorithms based on Markov chains are used widely in the sciences. The structure of these models makes them amenable to analysis via linear algebra, and we will explore this beautiful theory during our first unit.
2. **Martingales**. Martingales are a class of random process in which the change at each step is unbiased. Examples include gambling and fair games, markets, empirical losses in machine learning, and more. Though Martingales have very little structure, we will be able to make strong statements about their properties and behavior. We will use these properties to study concentration of many random processes, and to better understand Markov Chains.

3. **Poisson Processes**. A Poisson Point Process is a continuous-time (or -space) model for the occurrence of independent events; Poisson processes arise in physics, biology, human behavior, and more. The structure of these processes yields an elegant theory, which we will develop and then use to study queuing theory and (in the next unit) continuous-time Markov Chains.
4. **Continuous-time Markov chains**. Continuous-time Markov chains are Markov chains where steps occur in "continuous time" rather than at discrete time intervals. This allows us to model a richer class of processes, generalizing many of the tools from discrete-time Markov chains. We'll introduce the basic concepts and see applications in queuing theory.
5. **Special topics: interacting particles & statistical physics**. We'll see some of the ideas we have studied in the context of statistical physics, seeing how the machinery that we have built up allows us to predict phase transitions and understand "physical characteristics" of the Ising model, a simple model of interacting particles in a material.

A fine-grained course schedule may be found on the website.

3 Materials & Resources

Course website. The course website is www.tselilschramm.org/stochproc/stats217-winter23.html. There you will find the course schedule, a list of texts and resources, and additional relevant readings.

Cavnas. The course will be recorded, and recording will be available on Canvas. We'll also use Canvas to post Zoom links for remote lectures, if necessary.

Gradescope. Homeworks and solutions will be made available on Gradescope, and you will turn in your homework solutions on Gradescope. You will receive a Gradescope invitation in the first week of the quarter.

4 Coursework & Evaluation

Homework (50%) We will have five homework assignments in total, each weighted equally. Your lowest homework score will be dropped. Collaboration is encouraged, but you must independently write your own solutions.

Problem Set Reflections (20%) Within a week of submitting your problem set, you will self-grade your problem set, and submit, for each question, a brief discussion of whether/how your solutions differ from the answer key. The goal here (as the name suggests) is for you to reflect on your solution, so that doing the problem set is a more useful experience.

If you believe your answer to be correct, you may just mark your solution with a check mark or similar. Note however that if you mark an incorrect solution correct, you will not receive points for the reflection.

In order to make grading easier for the course staff, you must **submit your reflection as well as your original solution (even if your original answer was correct)**; your reflection should appear below (or next to) your original answer.

Final Exam (30%) There will be a final exam during finals week, details forthcoming. We will provide a practice final exam during the week of March 7 to help you prepare.

Course grades Students whose grade falls in the range

- 85% – 100% will receive between A- and A+.
- 70% – 84% will receive between B- and B+.
- 55% – 69% will receive between C- and C+.
- 45% – 54% will receive between D- and D+.
- 0% – 44% will receive a NP.

To receive credit for the class with CR/NC grading, you must obtain a minimum of C-.

5 Policies

The Honor Code. It is expected that you and I will follow Stanford's Honor Code in all matters relating to this course. You are encouraged to meet and exchange ideas with your classmates while studying and working on homework assignments, but you are individually responsible for your own work and for understanding the material. You are not permitted to copy or otherwise reference another student's homework or computer code.

Late Work Policy. Late work will not be accepted. To allow you some flexibility, your lowest homework and reflection scores will be dropped.

Accommodations. I am happy to provide accommodations, understanding that they may be necessary for student success. 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). Students should contact the OAE as soon as possible since timely notice is needed to coordinate accommodations.

Course Privacy Statement. As noted in the University's policy on recording and broadcasting courses, students may not audio or video record class meetings without permission from the instructor (and guest speakers, when applicable).