Syllabus

1 Logistics

- Instructor: Tselil Schramm (tselil@stanford.edu).
- Course website: www.tselilschramm.org/statstheory/stats300b-winter24.html.
- Teaching Assistants: John Cherian and Tim Sudijono.
- Lectures: Mondays & Wednesdays, 11:30-12:50, 60-109.
- Office Hours: listed on Canvas.

The best way to contact me is by email. Please be sure to include "STATS 300B" in the subject line.

Prerequisites: Students will need a background in real analysis (MATH 171), probability theory (STATS 310A), and a solid command of linear algebra. Some finite-sample hypothesis testing theory (e.g. STATS 300A) is helpful, but not essential. 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 covers core topics in asymptotic statistics and high dimensional statistics. Our goal will be to understand (through a hard-line mathematical lens) properties of the outputs of learning algorithms (a.k.a. estimators) computed from a large quantity of (usually i.i.d.) data. We will explore the following questions: How do estimators behave (when we have access to ample i.i.d. data)? How can we quantify their success? What makes some estimators better than others? En route to answering these questions, we will introduce a variety of useful tools in concentration of measure and analysis.

Topics The course is roughly partitioned into the following units:

1. **Stochastic Convergence and Central Limit Theorems**. We begin with some technical tools that form the basis for asymptotic statistics: notions of stochastic convergence, the law of large numbers, the central limit theorem, and the delta method.
2. **Maximum Likelihood, M-estimators, and Testing**. In this unit we will study the limiting behavior of maximum likelihood estimators (MLE) and other M-estimators (solutions to optimization problems) on i.i.d. data, as well as basics of hypothesis testing. The punchline for estimators is roughly that, under some "niceness" conditions, the law of such estimators is normal in the limit as the number of samples goes to infinity. Topics include asymptotic normality, consistency, influence, confidence intervals, and a variety of hypothesis tests.
3. **Optimality Theory**. Next we discuss the efficiency of estimators and what makes some estimators preferable to others, building up to the optimality of the MLE. Topics include contiguity of probability measures, Le Cam's lemmas, quadratic mean differentiability, local asymptotic normality, and the local asymptotic minimax theorem.

4. **Non-Asymptotic Rates (High-Dimensional Statistics)**. We then go beyond the asymptotic setting, building a theory that incorporates the dependence of rates on dimension and other properties of the data. The emphasis shifts from determining the precise law of limiting statistics, to instead obtaining reasonable non-asymptotic confidence intervals. Topics include concentration inequalities, matrix concentration, uniform convergence, Rademacher complexity, VC dimension, Gaussian processes, chaining. The mathematical tools introduced here apply equally well in the more classical asymptotic theory.
5. **Computational Efficiency**. In many cases, the MLE is not computable by polynomial-time algorithms (unless P = NP). We will discuss a few of the algorithms/heuristics commonly employed to overcome this, and analyze (using our STATS 300B toolkit) the performance of these heuristics in some contexts where they provably succeed.

The course schedule may be found on the website.

3 Materials & Resources

Course website. The course website is www.tselilschramm.org/statstheory/stats300b-winter24.html. There you will find the course schedule, a list of texts and resources, and additional relevant readings.

Scribe notes on Overleaf. Each lecture will be scribed by a student, and the scribe notes will be available for all to edit on Overleaf. You will be added to the Overleaf project in the first week of the quarter.

Canvas. Homework solution keys, Tselil's handwritten lecture notes, and any other files will be made available on Canvas.

Gradescope. Homework assignments will be posted 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.

Campuswire. We will use Campuswire as an online class forum, where you may ask questions and discuss with your fellow students. You will receive a Campuswire invitation in the first week of the quarter.

4 Coursework & Evaluation

Homework (45%) We will have problem sets every week, each of equal weight. We will drop your two lowest homework scores. Collaboration is encouraged, but you are responsible for your own understanding and you must independently write your own solutions.

Problem Set Reflections (optional) Within a week of receiving your graded problem set, you may submit a brief discussion of whether/how your solutions differ from the answer key (an example will be posted on Canvas). If you demonstrate an understanding of your mistake, you may earn back up to 50% of any deducted points (for problems on which you made a good faith attempt the first time around).

Scribe Notes (10%) You will be responsible for the scribe notes of one lecture. Students will sign up (possibly in pairs) to scribe each lecture; students have 48 hours to type up the scribe notes.

The lecture notes will be available to edit by all on Overleaf; you're encouraged to add to and improve the notes for any week, regardless of whether it is the week you are responsible for. Overleaf allows us to see who made which edits; your contributions across all 19 lectures of the class will be counted.

Full credit will only be given to students who produced complete, useful notes! Some comments on this: In a lecture, the instructor usually only writes a subset of what they are communicating on the board. Good scribe notes are not just a latex transcript of what the instructor wrote; they fill in the gaps. For this reason, producing scribe notes requires a thorough understanding of the material, which often entails some additional reading of the supplementary texts.

Final Exam (45%) The final will be in-person during finals week. Students are expected to be present for the final.

Campuswire participation (up to 5% bonus) Students may receive up to 5% extra credit for meaningful participation on Campuswire. Meaningful participation constitutes the asking or answering of questions in a way that has a positive impact on the learning of other students in the class.

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. Instead, your lowest two homework 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).