# 36-771 : Martingales 1 (Concentration inequalities)
# 36-772 : Martingales 2 (Sequential analysis)

2 minis (Aug 28 to Oct 18, Oct 23 to Dec 6), 6 credits each
Fall 2018, Syllabus

August 25, 2018

## 1 Basic Course Information

**Instructor** Aaditya Ramdas, aramdas@stat.cmu.edu [Office hours: on demand]
**Teaching Assistants**: there are no TAs for this course.

**Time:** Tue, Thu 10:30-11:50am
**Location**: Wean 4625
**Exceptions**: There will likely be no class on Sep 13 (Thu), Nov 1 (Thu) due to instructor travel, Oct 16 due to university rules, and on Nov 20 (Tue), Nov 22 (Thu) due to Thanksgiving. Hence there will be 14 lectures in Mini 1, and 11 lectures in Mini 2 (Mini 1 is longer than Mini 2 by design). See the CMU academic calendar.

**Website** See http://www.stat.cmu.edu/∼aramdas/martingales18 for basic course material.
**Announcements** All announcements will be made on Canvas.

**Participants** This course is intended for advanced PhD students with strong mathematical background.
**Prerequisites** The first mini is a prerequisite for the second mini. There are no formal prerequisites for the first mini, but non-PhD students must email the instructor for permission to enroll in the course. Students are expected to have completed at least one intermediate statistics course (like 10-705 at CMU), and preferably an advanced statistics course. Students must be familiar with

1. basic concentration inequalities (such as Markov, Chebysheff, Hoeffding)
2. basics of martingales (filtrations, stopping times, Doob's maximal inequalities and optional stopping theorems)
3. basics of testing and estimation (probability ratio tests, risk, type-1 error, power)
4. basics of convex analysis (strict convexity, Legendre-Fenchel transform)

**Textbook** There is no single textbook we will follow. Some related textbooks include:

- Large deviations techniques and applications, by Dembo and Zeitouni
- Stopped random walks: limit theorems and applications, by Gut
- Self-normalized processes, by de la Peña, Lai and Shao
- Concentration inequalities for sums and martingales, by Bercu, Delyon and Rio
- Concentration inequalities: a nonasymptotic theory of independence, by Boucheron, Lugosi and Massart

The presentation in Mini 1 will broadly follow a recent paper: Exponential line-crossing inequalities. Other relevant papers will be distributed as necessary.

## 2 Course Description

**36-771: Martingales 1 (Concentration inequalities)** Martingales are a central topic in statistics, but are even more relevant today due to modern applications to sequential learning and decision making problems. The first mini will present a unified derivation of a wide-variety of new and old concentration inequalities for martingales. We will prove inequalities for scalars and matrices, that hold under a wide variety of nonparametric assumptions.

Note to those who have already taken Advanced Probability or Advanced Statistics, you will also learn about (a) self-normalized exponential concentration inequalities for heavy-tailed distributions like those with only two moments, and even for those with no moments (like the Cauchy), (b) concentration of continuous time processes like Brownian motions (and possibly, Poisson processes, Levy processes), (c) concentration of martingales in smooth Banach spaces (useful for vector and matrix concentration).

**36-772: Martingales 2 (Sequential analysis)** The second mini will focus on deriving guarantees for a variety of important problems in sequential analysis using the tools developed in the first mini, as well as new tools such as uniform nonasymptotic versions of the law of the iterated logarithm for scalars and matrices. Applications include sequential analogs of the t-test that are valid without a Gaussian assumption, best-arm identification in multi-armed bandits, average treatment effect estimation in sequential clinical trials, sequential covariance matrix estimation, and other such problems.

## 3 Graded Components

**Mini 1** The grade will be based on one long homework (40%) whose questions will be progressively released, and a course project (10% for proposal, 50% for final report). Homework 1 will be due on Oct 5 (Fri). The project proposal (approximately one page) is due on Sep 13 (Thu), and the final report (approximately five pages) due on Oct 16 (Tue, when class is cancelled due to university policy). The course summary, reflection and open problems discussion will occur on the last day of the mini, Oct 18.

**Mini 2** 40% of the grade will be based on one long homework whose questions will be progressively released. For the rest, students will preferably develop a significant extension of the course project in Mini 1 (50% for final report, 10% for a short in-class presentation). Homework 2 will be due on Nov 15 (before Thanksgiving week). The in-class presentations will occur on Dec 4 (Tue). The project final report (approximately ten pages) will be due on Dec 5 (Wed). The course summary, reflection and open problems discussion will occur on the last day of classes, Dec 6.

**Projects** There are a wide variety of options available for course projects. Examples include:

- (Low risk) You can survey an area of the literature (covered in a textbook, or a set of advanced papers) that is related to the course, and is complementary to what is covered in class.
- (Medium-low risk) You can create a set of graphs, plots, or interactive figures, which allow the user to visualize several of the concentration inequalities and/or applications covered in the course. For inspiration, check out distill.pub, and specifically, a paper on why momentum works.
- (Medium-high risk) You can apply the contents of the class to your own research problem, for example by improving the guarantees you had achieved by using older tools, or by extending the analysis of your problem to hold for new settings.
- (High risk) If you are mathematically very mature, and want to work on a new research problem in this area from scratch, talk to the instructor privately in person.

Other ideas for course projects are also welcome. Grades will ultimately be awarded based on the instructor's judgment of the amount of work completed in the project, with some amount of subjective discounting for the risk of the project taken up.

## 4 Learning Objectives

**Mini 1** Upon successful completion of the first mini, the student will be able to

- Identify mathematical expressions that correspond to typical martingales and supermartingales.
- Quote Ville's maximal inequality for supermartingales, and explain it using a figure.
- Verify that "the central assumption" of the course is satisfied under various nonparametric conditions.
- State and apply "the mother of all exponential concentration inequalities", the main theorem of the course.
- Map an "exponential line-crossing inequality" onto a traditional concentration inequality.
- Assess which of two concentration inequalities is more general, using the A-B-C-D-E method.

**Mini 2** Upon successful completion of the second mini, the student will be able to

- Describe the difference between a confidence interval and a "confidence sequence".
- Explain how to derive an "exponential curve-crossing inequality" from a line-crossing inequality.
- Translate a nonasymptotic law-of-iterated-logarithm concentration inequality into a confidence sequence.
- Derive variants of the sequential probability ratio test using the supermartingale technique.
- Assess the pros and cons of the standard t-test, compared to the LIL-t-test.

## 5 Approximate Schedule (will change adaptively)

**Mini 1** (14 lectures)

- Intro: Concentration of measure, Doob's & Ville's inequalities, optional stopping, examples (1-2 lectures).
- The central assumption and sufficient conditions: sub-Gaussian, sub-exponential, sub-Poisson, etc (2 lectures).
- The main theorem: four equivalent statements (2 lectures).
- Generalizations of Hoeffding, Bennett, Bernstein, Freedman, etc (2 lectures).
- Heavy-tails, and self-normalized inequalities (2 lectures).
- Strengthening standard matrix concentration inequalities (2 lectures).
- Continuous-time processes and/or Banach-space concentration (1-2 lectures).
- Summary, reflection and open directions (1 lecture).

**Mini 2** (11 lectures)

- Intro: uniform, nonparametric, nonasymptotic confidence sequences, inadequacy of line-crossing (2 lectures).
- Stitching & mixing : from exponential line- to curve-crossing inequalities (2 lectures).
- Inverted stitching : numerically estimating curve-crossing probabilities (1 lecture).
- Application 1: a nonasymptotic, nonparametric, uniform sequential t-test (1 lecture).
- Application 2: multi-armed bandits, best-arm identification, adaptive hypothesis testing (1 lecture).
- Application 3: sequential covariance matrix estimation (1 lecture).
- Application 4: sequential average treatement effect estimation (1 lecture).
- Project presentations (1 lectures).
- Summary, reflection and open directions (1 lecture).

# 6 Course policies

## 6.1 Attendance

It is not compulsory, but recommended and expected. Every research study on this topic that I have read concludes that academic performance is negatively affected by not showing up to class. To encourage attendance, subtle hints for exam questions will be dropped from time to time.

## 6.2 Collaboration

Discussion of class material is heavily encouraged. Additionally,

- After submission of a homework, discussion of answers is encouraged.
- Before submission of a homework, reasonable verbal discussion of homeworks is allowed. An example of unreasonable verbal discussion: one person reciting formulae orally while another one writes them down. Written discussion (in any form) is permitted in groups smaller than 3 (or in rare exceptions 4) students.
- No matter what discussions have taken place, every homework and cheat sheet and mini-project and self-test (in its entirety) must be written up or coded up alone.

## 6.3 Academic Integrity

I have a zero tolerance policy for violation of class policies. If you are in any doubt whether a form of collaboration or obtaining solutions is permitted, please clarify it with me before proceeding.

- For each question on each homework, collaborators for that question must be acknowledged. Copying solutions from the internet is explicitly disallowed. You may search for material to help you understand a concept better, but be sure to create your own final solution. If you happen to use results from Wikipedia or textbooks, you must cite the source and are expected to completely understand the result you are citing. However, it is disallowed to copy solutions to exercises from elsewhere on the internet, like other courses or papers. When quoting text from a textbook, paper or website, use the `\begin{quote}` option in Latex.
- Any deviation from the rules will be dealt with according to the severity of the case. For example: evidence of written discussion in a larger group than 3-4 will result in points earned for that question becoming zero for all those relevant students; blindly copying one solution from someone else or online will result in the maximum points that can be earned for that homework becoming zero (maximum eligible grade becomes B); repeat occurrences will result in a failing grade for the course.
- In line with university policy, all instances of cheating/plagiarism will be reported to your academic advisor and the dean of student affairs. See the university policy on academic integrity.

## 6.4 Use of Mobile Devices and Laptops in Class

These are allowed but not encouraged. Learning research shows that unexpected noises or movement automatically divert and capture people's attention, meaning that you are affecting everyone's learning experience. For this reason, I ask you turn off your mobile devices and close your laptops during class. If you must use your laptop or mobile, make sure you are sitting at the back of the class.

## 6.5 Late Assignments

Every student is allowed a total of 2 late days per mini. Beyond that, the maximum earnable points for that assignment will drop by 20% per day.

# 7 Additional information

## 7.1 Accommodations for Students with Disabilities

If you have a disability and are registered with the Office of Disability Resources, I encourage you to use their online system to notify me of your accommodations and discuss your needs with me as early in the semester as possible. I will work with you to ensure that accommodations are provided as appropriate. If you suspect that you may have a disability and would benefit from accommodations but are not yet registered with the Office of Disability Resources, I encourage you to contact them at access@andrew.cmu.edu.

## 7.2 Statement of Support for Students' Health & Well-being

Take care of yourself. Do your best to maintain a healthy lifestyle this semester by eating well, exercising, avoiding drugs and alcohol, getting enough sleep and taking some time to relax. This will help you achieve your goals and cope with stress.

If you or anyone you know experiences any academic stress, difficult life events, or feelings like anxiety or depression, we strongly encourage you to seek support. Counseling and Psychological Services (CaPS) is here to help: call 412-268-2922 and visit http://www.cmu.edu/counseling/. Consider reaching out to a friend, faculty or family member you trust for help getting connected to the support that can help.

If you or someone you know is feeling suicidal or in danger of self-harm, call someone immediately, day or night (CaPS: 412-268-2922, Resolve Crisis Network: 888-796-8226). If the situation is life threatening, call the police (On-campus CMU Police: 412-268-2323, Off-campus Police: 911).