

Information Sheet

Lecturers

Polina Golland	Gregory W. Wornell
Office: 32-D470	Office: 36-677
Tel: x3-8005	Tel: x3-2297
E-mail: polina@csail.mit.edu	E-mail: gww@mit.edu

Administrative Assistant

Tricia O'Donnell
Office: 36-677
Tel: x3-2297
E-mail: tricia@mit.edu

Teaching Assistants

Igor Spasojevic	E-mail: igorspas@mit.edu
Chandler B Squires	E-mail: csquires@mit.edu
Matthew Staib	E-mail: mstaib@mit.edu
Jennifer Susan Tang	E-mail: jstang@mit.edu

Lectures: Tuesday and Thursday, 9:30am–11am, Room 32-123
Recitations: Friday, 1–2pm or 2–3pm, Room 4-163
Office Hours: Days, times and locations posted on stellar site.

Welcome to 6.437!

This course offers an immersive graduate-level introduction to the principles of statistical inference, with an emphasis on information theoretic perspectives. As such, it is a core graduate subject for students in the relevant subfields both within and beyond Course 6 (EECS). The material in this course constitutes a common foundation for work in, for example, machine learning and data science, artificial intelligence, signal processing, computer vision and computational imaging, communication, bioinformatics and neuroscience, finance and data analytics, and network science.

It is worth stressing that 6.437 is an *introductory* graduate subject: it is *not* an advanced graduate subject for students who have already mastered statistics, estimation and decision theory, and information theory, and now want to understand such material at an even more sophisticated level. Nevertheless, by its structure,

6.437 will ultimately reveal connections among these fields, and provide background for more advanced treatments.

Ultimately, the course is about teaching you contemporary approaches to—and perspectives on—problems of statistical inference. The development of the material that forms the basis for this subject has historically been strongly driven by applications, and continues to be. However, our focus in the course will not be on these applications—which form the basis for entire courses of their own—but rather on the common problem-solving frameworks that they share. Nevertheless, we will cite various relevant applications as we develop the material and sometimes extract simplified illustrative examples from these contexts.

Note that the course has both lectures and recitations, which are designed to complement each other. Recitations begin the first week of classes. There are two possible recitation times to choose from, as indicated above. Select to attend whichever suits your schedule best. In addition, there are office hours scheduled throughout the week, as indicated on the 6.437 web site. You are welcome and encouraged to come to any and all of them you think might be helpful to you in clarifying your understanding of the material.

Prerequisites

The official prerequisite is 6.008, 6.041/6.431, 18.05, 18.600, or 6.436. The effective prerequisite is *fluency* with basic quantitative probabilistic reasoning and analysis, together with the kind of mathematical maturity that often comes from taking at least one higher level undergraduate subject that has a significant mathematical component. As such, a student having had 6.436 would be sufficiently well prepared, while a student having only had 6.008, 6.041/6.431, 18.05, or 18.600 and no subsequent subjects of a strong mathematical flavor would likely need additional preparation. As an example, having had one of these subjects together with an introductory subject in analysis (e.g., 18.100), would be sufficient, but not necessary, preparation.

When in doubt, students whose undergraduate degrees are not from MIT should consult the staff to determine if they have had subjects that are effectively equivalent to the official prerequisites.

Reading

There is no existing text that matches the content of this somewhat unique subject and the style in which we teach it. However, we have been developing a set of detailed course notes, which we will distribute in parts as we go along. These notes are under perpetually active development, and as such are necessarily rough in places and contain bugs, which we will count on you to help us catch.

If you are interested in further reading, either to strengthen your background, reinforce some of the concepts from lecture, or to probe some topics in more detail,

you might want to take a look at the additional references on the course web site, which have been selected with some care. In particular, you'll find several papers containing a variety of useful insights, which are worth the effort to work through.

Problem Sets and Exercises

There will be 9 problem sets. Problem sets will be due in lecture (except the final problem set, which is never due). Problem sets must be handed in by the end of the class in which they are due. Problem set solutions will be available at the end of the due date's lecture.

While you should do all the assigned problems, only a randomly chosen subset will actually be graded. You will find some problems in the problem sets marked as "practice". These are not required, but you might find it helpful to work through them if you are looking for more practice working with the concepts introduced in class.

Don't be misled by the relatively few points assigned to homework grades in the final grade calculation! While the grade you get on your homework is only a minor component of your final grade, working through (and, yes, often struggling with at length!) the homework is a crucial part of the learning process and will invariably have a major impact on your understanding of the material. Some of the problem sets involve a PYTHON component, to help you explore different aspects of the material.

In undertaking the problem sets, moderate collaboration in the form of joint problem solving with one or two classmates is permitted provided your writeup is your own.

To help you get started with the problem sets, there will be an accompanying exercises handout posted on the web site at the same time. These exercises are simple worked examples that illustrate some of the basic calculations that the problem sets will in turn build on. You're free to use these exercises any way you find helpful to you in preparing for the homework.

Project

In addition to the problem sets, the class will include a two-part PYTHON-based class project. Part I will be a guided exercise, while the continuation Part II will be a more open-ended challenge. The project is intended to be educational, interesting, and fun! The project will go out on Tuesday, April 23 and Part I will be due the following Tuesday. There will be no problem set issued during Part I of the project. Part II will be due on Friday, May 10. You may want to make a note of the project dates in your schedule now, to help you with planning your time and coordinating with other classes during the semester.

The project details will be announced closer to the time it goes out, but plan on an interesting and engaging experience with the material.

Exams

There will be two quizzes in the subject. The first quiz will be held Wednesday evening, March 13, 7-10pm. The second quiz will be held during finals week (the date and time for which will be determined by the Registrar). The quizzes will be designed to require 1.5 hours of effort, but we'll use the three hour format to minimize the effects of time pressure. The quizzes will both be *closed book*. You will be allowed to bring *two* 8.5 × 11-inch sheet of notes (both sides) to the Midterm Quiz, and *four* 8.5 × 11-inch sheets of notes to the Final Quiz.

Course Grade

The final grade in the course is based upon our best assessment of your understanding of the material during the semester. Roughly, the weights used in grade assignment will be:

Midterm Quiz	40%
Final Quiz	40%
Project	5%
Homework	15%

with an additional property that if you do better on the Final Quiz than the Midterm Quiz, and you have done all the problem sets and the project, then the Midterm Quiz will not count, i.e., the Midterm Quiz can only help you if you are doing all the homework. Also, the project must be completed to receive a grade in the subject.

As always, other factors such as contributions to the lecture discussion and other interactions can make a significant difference in the final grade.

Course Web Site, Discussion Forum, and Email

We will make announcements via email, and we will post various information and handouts on the course web site.

You should first make sure that you have an active Athena account (by visiting <http://ist.mit.edu/support/accounts> if necessary) as well as a personal certificate (by visiting <https://ca.mit.edu/ca/> if necessary). If you have problems or if you are not a regular MIT student, please contact one of the TAs for assistance.

The course web site is

<http://web.mit.edu/6.437>

You will need to have a valid certificate *and* be on the official course list to access the web site. If you have pre-registered for 6.437, this should already be set up; just double-check that you can access the web site (try to download a handout, for example). Otherwise, contact one of the TAs and they will add you to the list. If

you can access content on the web site, you should also be receiving all of the course announcements.

Additionally, we will use PIAZZA to post class announcements and as an on-line discussion forum. Through piazza you can post questions or comments (anonymously or otherwise) about any aspects of the material, which your peers and the staff can respond to. It has become a popular resource in the course for clarifying potential points of confusion in lecture notes, problem sets, recitation handouts, project descriptions, etc. You can access the piazza forum through the link provided on the 6.437 web site. Once you enroll yourself, you can fully participate in the discussions.

Finally, if you have other kinds of questions during the term, you can always reach us by sending email to

`6.437-staff@mit.edu`

Tentative Syllabus and Schedule

	Date	Topic	Due	Out
T	2/5	L1: Introduction and overview		PS1
R	2/7	L2: Bayesian hypothesis testing and estimation		
T	2/12	L3: NonBayesian decision theory	PS1	PS2
R	2/14	L4: Minimax decision theory		
T	2/19	<i>Monday schedule – no class</i>		
R	2/21	L5: NonBayesian parameter estimation		
T	2/26	L6: Exponential families	PS2	PS3
R	2/28	L7: Sufficient statistics		
T	3/5	L8: The EM algorithm	PS3	PS4
R	3/7	L9: Inference as decision		
F	3/8	<i>Add Date</i>		
T	3/12	L10: Information geometry	PS4	PS5
W	3/13	(Evening) Midterm Quiz (through L9 and PS4)		
R	3/14	<i>no class</i>		
T	3/19	L11: Modeling as inference		
R	3/21	L12: Extensions to continuous parameters		
	3/25-3/29	<i>Spring Break</i>		
T	4/2	L13: Priors	PS5	PS6
R	4/4	L14: Alternating projections		
T	4/9	L15: Approximate inference via sampling	PS6	PS7
R	4/11	L16: Asymptotics: typical sequences, large deviations		
T	4/16	<i>Patriots Day vacation – no class</i>		
R	4/18	L17: Method of Types and Sanov’s Theorem		
T	4/23	L18: Asymptotics of hypothesis testing, estimation	PS7	Project
R	4/25	L19: Asymptotics of model capacity		
		<i>Drop Date</i>		
T	4/30	L20: Universal prediction and inference	Project, Part I	PS8
R	5/2	L21: Nested classes and model order selection		
T	5/7	L22: Learning high-dimensional models	PS8	PS9
R	5/9	L23: Universal feature characterizations		
F	5/10		Project, Part II	
T	5/14	L24: Logistic regression and neural networks		
R	5/16	L25: TBA		
	TBA	Final Quiz (through L25 and PS9)		