

Information Sheet

Lecturers

Prof. Polina Golland
Office: 32-D470
E-mail: polina@csail.mit.edu

Prof. Gregory Wornell
Office: 36-677
E-mail: gww@mit.edu

Teaching Assistants

Ganesh Ajjanagadde (gajjanag@mit.edu)
Woo Hyeok (David) Kang (wkang@mit.edu)
Julian Straub (jstraub@mit.edu)
Atulya Yellipedi (atulya@mit.edu)

Administrative Assistant

Tricia O'Donnell
Office: 36-677
Tel: 253-2297
E-mail: tricia@mit.edu

Lectures: Monday and Wednesday, 10–11 (32-155)
Recitations: Tuesday and Thursday, 11–12 (34-304), or 1–2 (36-155),
or 2–3 (36-155)
Office Hours: Days, times and locations posted on the web site

Welcome to 6.008 !

We live in a world in which we constantly need to extract information from data. This is the central problem of inference. And computationally efficient methods for such inference are enabling technologies for an enormous range of applications.

Example domains abound, and include search and retrieval, data mining, computer vision and imaging, voice recognition, communication and compression, natural language processing, robotics and navigation, computational biology and bioinformatics, medical diagnosis, distributed sensing and monitoring, and finance.

Many of the most successful inference algorithms arise out of probabilistic modeling and analysis. If you want to learn the fundamentals of this discipline and see some of what you can do with it, this subject is the place to start. Indeed, it will provide a solid foundation for more advanced subjects that build on this framework of reasoning. As such, the subject is targeted at (and likely to strongly appeal to) students both across and beyond Course 6 (EECS).

As with other foundational subjects, 6.008 has both lectures and recitations, which are designed to complement each other. Recitations begin the first week of classes. There are three possible recitation times from which you can choose, as indicated above. Select to attend whichever suits your schedule best; if strong imbalances result, we will make adjustments. In addition, there will be several scheduled staff office hours throughout the week. 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.

Although this subject has been taught twice in pilot form under a temporary number, this is the first time it is being offered under its permanent number, and thus 6.008 is a brand-new subject in our curriculum. As such, we will continue to engage in considerable experimentation with content and pedagogy. While this is an exciting and important part of course development, it also means that the subject will be rougher and more raw than a subject that has been offered many times. As a result, this inaugural offering of 6.008 will especially appeal to students who, as part of taking the subject, relish the role of beta-tester and are eager to contribute to the subject's formative development through their participation.

Prerequisites

The official prerequisite is 6.01, which includes a very brief (roughly two-week) introduction to probabilistic reasoning. But students who have had a comparable introduction in some other subject will also be suitably prepared. We emphasize that a more thorough introduction to probability, such as in the form of 6.041 or 18.05 or 18.440 is specifically *not* required, as we will develop the necessary foundation in probability as part of 6.008. However, we will assume you are comfortable with (and fluent in) basic mathematics to the level of, e.g., 18.02.

Degree Requirements

For EECS students, the requirements satisfied by 6.008 will change in the future as the department reviews the undergraduate curriculum. The requirements specified in this section are true for the Fall 2014 offering of 6.008.

EECS students in 6-2 program may use 6.008 as one of their EE or CS foundation subjects. All EECS students may use 6.008 as one of their math elective or free elective subjects.

All EECS students may petition to take 6.008 instead of 6.042 as one of their math elective subjects and use it as a prerequisite for more advanced subjects that require 6.042. Note that the petition must be filed and approved before the add date in Fall 2014. The goal is to verify that you either have sufficient background in constructing formal proofs (an essential component of 6.042 that more advanced subjects rely on) or that you are committed to acquiring such background on your own.

6.008 may also be used by MEng students as one of their restricted elective subjects. As always, no double counting is allowed.

Non-EECS students should consult their departments about what degree requirements 6.008 can satisfy for them.

If you have additional questions about degree requirements, please contact the lecturers. We will be happy to meet with you and discuss.

Reading

There is no existing text that matches the content of this new subject and the style in which we teach it. However, we have been actively developing course notes for the class, and will distribute them in parts as we go along. In this development stage, the notes will be necessarily rough in places and contain bugs, which we will count on you to help us catch.

In addition, you will find the following text a useful resource for the subject:

D. P. Bertsekas and J. N. Tsitsiklis, *Introduction to Probability*, 2nd ed., Athena Scientific, 2008.

If you'd like your own copy, which you will find handy, it may be purchased through, e.g., Amazon. We will also place this book on reserve at the MIT libraries (Barker). Other possibly useful readings will be posted on the course web site as appropriate.

We should emphasize that the course notes and text do *not* replace the lectures and recitations. The notes will be necessarily incomplete, and most of the discussion, interpretation, and insights will take place in class. Conversely, lectures and recitations are not a substitute for the notes and text. Indeed, we will defer many details of our development to the reading materials. Thus, you should view class and the written materials as strongly complementary.

Optional Kick-Start Sessions

We will hold two optional evening sessions on Wednesday, September 10, and Wednesday, September 17, 7-9pm. The goal of these sessions is to provide a bit more help getting started with probability and/or python for those who feel it would be beneficial to them. Each session will start with one hour on probability, followed by one hour on python. The format of the sessions will be an interactive tutorial led by a member of the staff.

Problem Sets and Computational Labs

There will be 11 problem sets (1-11) and 4 computational labs (I-IV). The problem sets emphasize analysis aspects of the material, while the computational labs involve a programming component, based on python, to help you explore computational aspects of the material.

Problem sets and computational labs will be concurrent (and jointly designed), so that at any time you will be working on one problem set and one computational lab. It will be important to manage your time carefully, and allow enough time for them.

Problem sets will be due by 10am on Wednesday, in lecture. Computational labs will be due by 5pm on Fridays and are to be submitted electronically. Solutions will be available on the course web site immediately after the due dates and times, and

thus we cannot accept late submissions for grading. To help with your planning, due dates are listed on the schedule at the end of this handout.

While you should do all the assigned problems and computational exercises, only a subset will actually be graded. Also, 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 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!) the homework is a crucial part of the learning process and will invariably have a major impact on your understanding of the material.

In undertaking the problem sets and computational labs, moderate collaboration in the form of joint problem solving with one or two classmates is permitted provided your writeup is your own and that you identify your collaborators in your writeup.

Exams

There will be two evening quizzes and a final exam in the subject. Dates for the quizzes are Tuesday, October 7, 7:30-9:30pm in 1-190 (conflict date: October 8, 9-11am), and Wednesday, November 5, 7:30-9:30pm in 32-155 (conflict date: November 6, 9-11am). Please let us know at least a month in advance if you have a conflict with the quiz time and need to make use of the conflict date. Note also that because of the evening quizzes, the corresponding Wednesday, October 8 and Wednesday, November 5 lectures in the schedule are canceled.

The final exam will be during Finals Week, the exact date of which we will announce when scheduled by the Registrar.

The quizzes and final exam will all be *closed book*. You will be allowed to bring *one* 8.5×11 -inch sheet of notes (both sides) to Quiz#1, *two* 8.5×11 -inch sheets of notes (both sides) to Quiz#2, and *three* 8.5×11 -inch sheets of notes (both sides) to the Final Exam.

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:

Quiz #1	20%
Quiz #2	20%
Final Exam	40%
Problem Sets	10%
Computational Labs	8%
Classroom participation	2%

with the additional property that if you do better on the Final Exam than either quiz, and you have done all the homework (problem sets and computational labs), then the corresponding quiz will not count, i.e., the quizzes can only help you if you are doing all the homework.

In addition, as always, other factors such as contributions to the discussion in class and other interactions can make a difference in the final grade.

Course Web Site and Email

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

You should first make sure that you have an active Athena account (by visiting <http://web.mit.edu/accounts/> if necessary) as well as a personal certificate (by visiting <http://web.mit.edu/ist/topics/certificates/> 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/6.008>

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 6.008, 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.

The student email list is

6.008-students@mit.edu

and will be kept in sync with the web site access list. If you can access content on the web site, you should also be receiving all of the course announcements.

If you have any questions during the term, you can reach us by sending email to

6.008-staff@mit.edu

Tentative Syllabus and Schedule

Date	Topic	HW out	HW due
W 9/3	L1: Introduction and overview	1	
M 9/8	L2: Discrete random variables, measures of randomness		
W 9/10	L3: Joint distributions, marginals, conditionals, Bayes	2	1
M 9/15	L4: Measures of randomness for joint distributions		
W 9/17	L5: Marginal and conditional independence structure	3	2
F 9/19		I	
M 9/22	L6: Decision-making, most probable configurations, MAP rule		
W 9/24	L7: Graphical models, message-passing, hidden Markov models	4	3
M 9/29	L8: HMM marginalization: forward-backward algorithm		
W 10/1	L9: HMM most probable configuration: Viterbi algorithm	5	4
F 10/3	<i>Add Date</i>	II	I
M 10/6	L10: Parameter estimation, Maximum Likelihood method		
T 10/7	Quiz 1 (through L7, PS4, and CL-I) no class		
W 10/8	no class		
M 10/13	<i>Columbus Day — no class</i>		
W 10/15	L11: Model learning; Naïve Bayes Models and HMMs	6	5
M 10/20	L12: Learning structure in graphical models		
W 10/22	L13: Markov/Chebyshev bounds, law of large numbers	7	6
F 10/24		III	II
M 10/27	L14: Typical sets, compression and hashing		
W 10/29	L15: Joint typicality, classification and communication	8	7
M 11/3	L16: Atypical sequences, large deviations and max-entropy		
W 11/5	Quiz 2 (through L13, PS7, and CL-II) no class		
M 11/10	<i>Veterans Day — no class</i>		
W 11/12	L17: Markov chains and random walks	9	8
M 11/17	L18: Sampling and approximate inference		
W 11/19	L19: Markov chain Monte Carlo and Gibbs sampling	10	9
F 11/21	<i>Drop Date</i>	IV	III
M 11/24	L20: Continuous random variables		
W 11/26	<i>Pre-Thanksgiving — no class</i>		
M 12/1	L21: Joint PDFs; continuous inference		
W 12/3	L22: Jointly Gaussian random variables, innovations	11 [†]	10
F 12/5			IV
M 12/8	L23: Gaussian inference and modeling; linear inference		
W 12/10	L24: Central limit theorem, Laplace approximations		
TBA	Final Exam (<i>Finals Week, 12/15-12/19</i>)		

[†]Not to be handed in.