Massachusetts Institute of Technology
Department of Electrical Engineering and Computer Science
6.S080 Introduction to Inference
Fall 2012

Information Sheet

Lecturers/In-Charge

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

Prof. Lizhong Zheng
Office: 36-660
E-mail: lizhong@mit.edu

Prof. Costis Daskalakis
Office: 32-G694
E-mail: costis@csail.mit.edu

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

Teaching Assistants

George Chen (georgehc@mit.edu)
Gauri Joshi (gauri@mit.edu)
Ramesh Sridharan (rameshvs@mit.edu)

Administrative Assistant

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

Lectures:	Monday and Wednesday, 10am–11am, Room 56-154
Recitations:	Tuesday and Thursday, 1-2pm or 3-4pm, Room 34-303
Office Hours:	Days, times and locations posted on web site

Welcome to 6.S080 !

We live in a world in which we constantly need to extract information from data. This is the central problem of inference. And computational 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 lan- guage 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 targetted at (and likely to strongly appeal to) students both across and beyond Course 6 (EECS). And students in 6-2 may even select 6.S080 as one of either their EE or CS foundation subjects.

As with other core subjects, 6.S080 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; 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.

As a brand-new subject in our curriculum, we will be engaged 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 offering of 6.S080 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 *not* required, as we will develop the necessary foundation in probability as part of 6.S080. However, we will assume you are comfortable with (and fluent in) basic mathematics to the level of, e.g., 18.02.

Reading

There is no existing text that matches the content of this new subject and the style in which we teach it. However, we will be actively developing course notes for the class, and will distribute them in parts as we go along. In this development stage, the notes will necessarily be 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 very useful resource for the subject, and recommend your purchasing a copy (for example, through Amazon). We will also place this book on reserve at the MIT libraries (Barker).

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

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 them. Thus, you should view class and the written materials as strongly complementary.

Problem Sets

There will be approximately 11 problem sets. Problem sets will be due by 12pm (noon) on Tuesdays. Your problems are to submitted to the prominent mail slot to immediately to the right of the door to 36-413 (RLE Headquarters). This slot is labeled "Headquarters Mail Drop-off," but don't be misled by the label; we have exclusive use of this slot for the term.

Problem set solutions (and the subsequent problem set) will be available on the course web site immediately thereafter, and thus we cannot accept late problem sets for grading. To help with your planning, problem set due dates are listed on the schedule at the end of this handout.

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!) 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 will involve a MATLAB component, to help you explore computational 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 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 Wednesday, October 3, 7:30-9:30pm, and Wednesday, November 7, 7:30-9:30pm. The final exam will be during Final's 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 to the Final Exam. The location of the quizzes will be announced when available. On the dates we have the quizzes, the lecture that would ordinarily take place is cancelled.

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 #3	40%
Homework	20%

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

In addition, as always, other factors such as contributions to the lecture discussion 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 various 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.s080`

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.S080, 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.s080-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.s080-staff@mit.edu`

Tentative Syllabus and Schedule

	Date	Topic
W	9/5	L1: Introduction and overview
M	9/10	L2: Discrete random variables, expectations, divergence, entropy
T	9/11	**Problem Set 1 issued**
W	9/12	L3: Joint distributions, marginals, conditionals, chain rule, Bayes
M	9/17	L4: Marginal and conditional independence, Markovianity, mutual info.
T	9/18	**Problem Set 1 due, Problem Set 2 issued**
W	9/19	L5: Decision-making, most probable configurations, MAP rule
M	9/24	L6: Graphical models, message-passing, hidden Markov models
T	9/25	**Problem Set 2 due, Problem Set 3 issued**
W	9/26	L7: HMM marginalization: forward-backward algorithm
M	10/1	L8: HMM most probable configuration: Viterbi algorithm
T	10/2	**Problem Set 3 due, Problem Set 4 issued**
W	10/3	**Quiz 1** (through L7 and PS3)
F	10/5	*Add Date*
M	10/8	*Columbus Day – no class*
W	10/10	L9: Illustrative Markov Chain and HMM applications
M	10/15	L10: Parameter estimation, Maximum Likelihood method
T	10/16	**Problem Set 4 due, Problem Set 5 issued**
W	10/17	L11: Model learning, Naïve Bayes Models and HMMs
M	10/22	L12: Markov/Chebyshev bounds, asymptotics, law of large numbers
T	10/23	**Problem Set 5 due, Problem Set 6 issued**
W	10/24	L13: Typical sets, entropy, compression and hashing
M	10/29	L14: *Hurricane Sandy – no class*
T	10/30	**Problem Set 6 due, Problem Set 7 issued**
W	10/31	L15: Joint typicality, mutual information, classification and communication
M	11/5	L16: Atypical Sequences, divergence, and large deviations
T	11/6	**Problem Set 7 due, Problem Set 8 issued**
W	11/7	**Quiz 2** (through L15 and PS7)
M	11/12	*Veteran's Day (observed) – no class*
W	11/14	L17: Markov chains and random walks, steady-state and mixing time
M	11/19	L18: Sampling and approximate inference, sample generation principles
T	11/20	**Problem Set 8 due, Problem Set 9 issued**
W	11/21	L19: Importance and rejection sampling
		Drop Date
M	11/26	L20: Markov chain Monte Carlo (MCMC) and Gibbs sampling
W	11/28	L21: Continuous random variables
M	12/3	L22: Joint, marginal, and conditional PDFs; continuous inference
T	12/4	**Problem Set 9 due, Problem Set 10 issued**
W	12/5	L23: Jointly Gaussian random variables, innovations representation
M	12/10	L24: Gaussian inference and modeling; linear inference
W	12/12	L25: Central limit theorem, maximum entropy, Laplace approximations
M	12/17	**Final Exam** (1:30pm-4:30pm, 37-212)