Massachusetts Institute of Technology
Department of Electrical Engineering and Computer Science
6.s080 Introduction to Inference
Spring 2014

Information Sheet

Lecturers

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 H Chen (`georgehc@mit.edu`)
Chelsea B Finn (`cfinn@mit.edu`)
Aldo Pacchiano Camacho (`pacchian@mit.edu`)
Oren Rippel (`rippel@mit.edu`)

Lectures:	Monday and Wednesday, 10–11 (32-155)
Recitations:	Tuesday and Thursday, 11–12 (5-134), or 12–1 (5-134), or 1–2 (13-4101), or 2–3 (13-4101)
Office Hours:	Days, times and locations posted on the 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 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 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 four 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.

Degree Requirements

The requirements satisfied by 6.s080 will change in the future as the department reviews the undergraduate curriculum. The requirements specified in this section are true for the Spring 2014 offering of 6.s080 only.

EECS students in 6-2 program may use 6.s080 as one of their EE or CS foundation subjects. All EECS students may use 6.s080 as one of their math elective or free elective subjects.

All EECS students may petition to take 6.s080 instead of 6.042 as one of their math elective subjects and to 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 Spring 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.s080 may also be used by MEng students as one of their restricted elective subjects. As usual, no double counting is allowed.

If you have additional questions about the 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 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 reading materials. Thus, you should view class and the written materials as strongly complementary.

Problem Sets

There will be 11 problem sets. Problem sets will be due by 10am on Wednesday, in lecture. Problem set solutions 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 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 programming component, to help you explore computational aspects of the material. We will use python for the programming exercises.

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, March 12, 7:30-9:30pm (conflict date: March 13, 8-10am), and Wednesday, April 16, 7:30-9:30pm (conflict date: April 17, 8-10am). On the dates we have the quizzes, the lecture that would ordinarily take place is canceled. Please let us know at least a month in advance if you have a conflict with the quiz time and plan to take advantage of the conflict date.

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 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	25%
Final Exam	40%
Homework	15%

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 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.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`

Syllabus and Schedule

	Date	Topic	PS out	PS due
W	2/5	L1: Introduction and overview	1	
M	2/10	L2: Discrete random variables, measures of randomness		
W	2/12	L3: Joint distributions, marginals, conditionals, Bayes	2	1
M	2/17	*Presidents' Day – no class*		
T	2/18	L4: Measures of randomness for joint distributions		
W	2/19	L5: Marginal and conditional independence	3	2
M	2/24	L6: Decision-making, most probable configurations, MAP rule		
W	2/26	L7: Graphical models, message-passing, hidden Markov models	4	3
M	3/3	L8: HMM marginalization: forward-backward algorithm		
W	3/5	L9: HMM most probable configuration: Viterbi algorithm	5	4
F	3/7	*Add Date*		
M	3/10	L10: Parameter estimation, Maximum Likelihood method		
W	3/12	**Quiz 1** (through L7 and PS4) no class		
M	3/17	L11: Model learning, Naïve Bayes Models and HMMs		
W	3/19	L12: Markov/Chebyshev bounds, law of large numbers	6	5
	3/24-28	*Spring break*		
M	3/31	L13: Typical sets, entropy, compression and hashing		
W	4/2	L14: Joint typicality	7	6
M	4/7	L15: Atypical Sequences, divergence, and large deviations		
W	4/9	L16: Markov chains and random walks	8	7
M	4/14	L17: Sampling and approximate inference		
W	4/16	**Quiz 2** (through L14 and PS7) no class		
	4/21-22	*Patriots Day – no class*		
W	4/23	L18: Importance and rejection sampling	9	8
R	4/24	*Drop Date*		
M	4/28	L19: Markov chain Monte Carlo and Gibbs sampling		
W	4/30	L20: Continuous random variables	10	9
M	5/5	L21: Joint PDFs; continuous inference		
W	5/7	L22: Jointly Gaussian random variables, innovations	11[†]	10
M	5/12	L23: Gaussian inference and modeling; linear inference		
W	5/14	L24: Central limit theorem, Laplace approximations		
	TBA	**Final Exam** (*Finals Week*)		

[†]Not to be handed in.