Massachusetts Institute of Technology
Department of Electrical Engineering and Computer Science
6.008 Introduction to Inference
Fall 2020

Information Sheet

Lecturers
Prof. Gregory W. Wornell
gww@mit.edu

Prof. Lizhong Zheng
lizhong@mit.edu

Teaching Assistants
Varkey Alumootil (varkey@mit.edu)
Sule Kahraman (sulek@mit.edu)
Jennifer Tang (jstang@mit.edu)

Lab Assistants
Tanner Bonner (tbonner@mit.edu)
Giray Kuru (ngkuru@mit.edu)
Jeremy Ma (jma22@mit.edu)
Danielle White (dmwhite@mit.edu)

Administrative Assistant
Tricia O'Donnell (tricia@mit.edu)

Lectures:	Monday and Wednesday, 10-11am
Recitations:	Tuesday and Thursday, 12-1pm or 1-2pm or 2-3pm (for the first week of classes)
Lab Hours:	Days, times and locations posted on web site.
Office Hours:	Days, times and locations posted on web site.

Note: All times are for Eastern Time Zone (ET), and all meetings virtual.

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 artificial intelligence and machine learning, search and retrieval, data mining, computer vision, computational imaging, voice recognition, natural language processing, robotics and navigation, computational biology and bioinformatics, medical diagnosis, distributed sensing and monitoring, communication and compression, economics and finance, social science and networking, and more.

Most of the most successful inference algorithms in use today arise out of probabilistic modeling and analysis. If you want to learn the fundamentals of this discipline and see some of interesting (and often surprising) things 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 students both within and beyond Course 6 (EECS).

Lectures, Recitations, Lab Sessions, and Office Hours

As with other core subjects, 6.008 has lectures and recitations, which are designed to complement each other. Recitations begin on Thursday of the first week of classes, i.e., the day after the first lecture. There are two possible recitation times from which you can choose. Select to attend whichever suits your schedule best; if strong imbalances result, we will make adjustments.

6.008 also has scheduled lab hours that while optional are highly recommended and will help you with the laboratory component of the class. For lab hours, we've set up a simple signup sheet (which you will access by clicking Materials tab in the nativation bar of the 6.008 web site and going to the "Signup Sheets" Section) where you can get yourself into the queue for help from an LA. Lab hours start the second week of classes.

In addition, there will be several scheduled staff office hours throughout the week, the details of which are provided on the class website. You are welcome and encouraged to come to any and all of them you think might be helpful to you in digesting the material.

For TA office hours, we've also set up a simple signup sheet (accessed by clicking the Materials tab in the nativation bar of the 6.008 web site and going to the "Signup Sheets" Section). Use this sheet to get yourself into the queue for help from a TA, and additionally specify what you have questions about, if you like. For example, you could list a particular part of a problem on the week's homework. We'll use this to help organize the office hours more efficiently, and group together students with similar questions. But we know that often you may have more general questions, and/or you may not even be able to quite describe what you are confused about, and that's fine! Just come and we'll help you. Office hours also start the second week of classes.

Lecture, recitations, lab sessions, and office hours will all be conducted on-line via Zoom this semester, due to the pandemic. Additional details are included later in this document. Video recordings of lecture, and of parts recitation, will be posted on the class web site. Additionally, a transcript of "virtual blackboard" content generated in each lecture will be posted for the class in real time, allowing you to refer back to material (asynchronously) as you are watching the lecture. You will access this via a dynamically updated link on Stellar, which you will access by clicking the Materials tab in the nativation bar of the 6.008 web site and going to the "Live Lecture Notes" Section.

Prerequisites

The official prerequisite is calculus. Effectively, we are assuming you are comfortable with (and fluent in) basic mathematical analysis to the level of, e.g., 18.02. In addition, we assume you are familiar with the basics of the Python programming language, but we will provide an (optional) tutorial in recitation on Tuesday, September

8 to help you come up to speed on the aspects of Python that will be needed in the laboratory component of the subject. Finally, we emphasize that no prior exposure to probability is required, as we develop the necessary foundations in probability as part of this class.

Degree Requirements

At the Institute level, 6.008 provides 12 units of Laboratory GIR credit. 6.008 also provides further credit towards degrees in a variety of departments. Consult the respective departmental web sites for details. Examples include the following:

- For students in 2-A (Engineering), 6.008 is a recommended introductory restricted elective.
- Students in 6-1 (Electrical Science and Engineering) can use 6.008 to meet their probability subject requirement, or as a departmental or unrestricted elective.
- Students in 6-2 (Electrical Engineering and Computer Science) can use 6.008 as one of their foundation subjects, and meets their probability and Level 1 EECS subject requirements. It can also serve as a departmental or unrestricted elective.
- Students in 6-3 (Computer Science and Engineering) can use 6.008 as a departmental or unrestricted elective.
- Students in 6-9 (Computation and Cognition) can use 6.008 as one of their core required subjects.
- Many students in 6-3, 6-7, and 6-14 have, in previous years, successfully petitioned to use 6.008 in place of the required subject 6.042, and used it as a prerequisite for more advanced subjects that require 6.042. If you plan to pursue this, you'll want to be sure you have (or will have) sufficient background in constructing formal proofs, which is an essential component of 6.042 that more advanced subjects rely on.
- Students in 11-6 (Urban Science) can use 6.008 to fulfill one of their computer science requirements.
- Students in 18-C (Mathematics with Computer Science) may use 6.008 as one of their restricted electives.

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

Reading

There are extensive and detailed course notes for the class, which we continue to actively develop. You will access these notes by clicking the Materials tab in the nativation bar of the 6.008 web site and going to the "Course Notes" Section.

The notes will be posted one section at a time over the course of the semester, with each installment posted in advance of its corresponding lecture(s). This section of the the web site also includes a "Course Notes Guide" document that is updated throughout the semester, and lists the specific sections of the reading that correspond to each lecture. As mentioned in the first installment of these notes, all sections marked with a "$*$" after the number are optional reading; read them as your interests dictate and your time permits.

Given their ongoing development, the notes will be invariably a bit rough in places and contain bugs, which we will be grateful to you for helping us catch. Other possibly useful readings will be posted on the course web site as appropriate.

We emphasize that the course notes do *not* replace the lectures and recitations. The notes will be necessarily incomplete, and many of the perspectives, interpretations, and insights will be developed in class. Conversely, lectures and recitations are not a substitute for the notes. 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.

Worked Exercises and Videos

To supplement the course notes, we have an evolving collection of worked exercises that demonstrate how to apply the concepts in simple settings. We will make such exercises available in conjunction with the problem sets. You may find these useful to go through after reading the relevant sections of the course notes as a way to help you get started with the homework. We emphasize, however, these posted exercises are entirely optional reading—use them if/however you think they will be helpful to you.

In addition to written materials, we have also been developing a repository of Khan Academy style videos to help you digest the reading and further explain concepts in the subject. These videos are optional viewing; they are another supplementary resource for the class.

Problem Sets and Computational Labs

There will be 10 problem sets (1-10) and 5 computational labs (I-V). The problem sets emphasize analytic aspects of the material, while the computational labs involve a programming component, based specifically on Python Version 3, 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 5pm on Wednesday, and submitted on-line via Gradescope (as discussed later in this document). Computational labs will be due by 5pm on their submission dates and are to be submitted on-line also via Gradescope. Solutions will be available on the course web site shortly after the due dates and times, and submissions after the solutions are posted will be recorded as completed, but not graded. To help with your planning, a comprehensive list of assignment and due dates is provided on the schedule at the end of this document.

While you should do all the assigned problems and computational lab components, only a subset will actually be graded. Also, you will find a set of "practice" problems in every problem set. These are not required, and you don't hand them in, but you may find it helpful to work through them if you are looking for more practice working with the concepts introduced in class.

It is important to emphasize that 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.

Many students find collaborating with one or two classmates on the problem sets and computational labs helpful, and we very much support and encourage this. Just remember that your writeup (and all code) needs to be your own, and you need to identify your collaborators in your writeup. You may find a slack/discord community, discussed later in this document, useful in finding classmates to work with, among other resources.

Exams

There will be a midterm quiz and a final exam in the subject, both of which you will take remotely via the class Gradescope site discussed later in this document. The midterm quiz will be held nominally Wednesday, October 21, 7:30-9:30pm, though you may choose to take this quiz in any two-hour window during the period Wednesday, October 21, 7:30pm to Thursday, October 22, 9:30pm. Note also that because of the midterm quiz, there is no lecture on Wednesday, October 21.

The final exam will be nominally three hours long, and take place remotely, during Finals Week. The final exam has been nominally scheduled for Thursday, December 17, 1:30-4:30pm, though you may choose to take this exam in any three-hour window during the period Wednesday, December 16, 1:30pm to Thursday, December 17, 4:30pm.

The quizzes and final exam will all be *closed book*. You will be allowed to have *two* 8.5×11-inch sheets of notes (both sides) for the Midterm Quiz, and *four* 8.5×11-inch sheets of notes (both sides) for 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:

Midterm Quiz	20%
Final Exam	30%
Problem Sets	20%
Computational Labs	25%
Class participation	5%

with the additional property that if you do better on the final exam than the midterm quiz, and you have done all the problem sets and attended at least 80% of the recitations, then the corresponding quiz will not count, i.e., the quiz can only help you if you are doing all the homework. Also, please note that completing all the computational labs is a subject requirement.

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

On-Line Platforms and Resources

We will make use of multiple on-line platforms and resources in 6.008.

Stellar, Athena, Certificates

We will make announcements via email, and we will post information and handouts on the 6.008 course web site, which resides on the Stellar platform:

http://stellar.mit.edu/S/course/6/fa20/6.008

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.

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

Zoom

To attend the lectures, recitions, lab hours, and office hours, you'll need to activate your MIT Zoom account here (click "Sign in"). Note that you will need the latest

version of Zoom installed, and that you will need MIT certificates installed on your machine to access our class Zoom sessions and content.

The Zoom Instructions document, accessible by the Zoom Info tab in the nativation bar of the 6.008 web site, describes the different ways we will use Zoom in lecture, recitation, lab hours, and office hours, and includes a very brief tutorial focusing on the relevant features. The Zoom Info tab also provides the different Zoom links we will use for each.

Gradescope

To submit homework, you will use Gradescope, which you can access via the Gradescope tab in the navigation bar of the 6.008 web site, logging in with MIT as your school credential and providing the entry code `943WYJ` as necessary.

Piazza

Additionally, we will use Piazza as an on-line discussion forum, which you can access (and enroll yourself in) via the Piazza Forum tab in the naviation bar of the 6.008 web site. Through Piazza you can post questions or comments (anonymously to your peers or otherwise) about any aspects of the material, which your peers and the staff can respond to. It is a popular resource in the course for clarifying potential points of confusion in class notes, problem sets, recitation handouts, computational labs, etc.

Slack/Discord

One of the most important aspects of the experience of taking a subject at MIT like 6.008 is the community of people you are taking the subject with. In normal times, this community forms and bonds naturally through side conversations and interactions among you in and outside of the classroom, as you navigate the semester together. And out of that shared experience emerge collaborations and friendships that endure long after the subject is over. Since due to the pandemic you won't be seeing each other in person in class, and since most of you won't even be on campus, we are eager to help you form this community using on-line tools. To this end, we will ask for an experienced volunteer or two from the class to set up and moderate a Slack workspace or Discord server for the class, which the staff (intructors, TAs, LAs, graders, and course assistant) will deliberately *not* be a part of. We encourage you to make the most of it, and while we will not monitor your Slack/Discord community, we remind you to be respectful and supportive of each other in your conversations.

Email

Finally, if you have administrative questions during the term, you can send email to

`6.008-admin@mit.edu`

We will be glad to help you, but please be sure that your question isn't already answered in this handout before you email us!

Tentative Syllabus and Schedule

Date		Topic	HW out	HW due
W	9/2	L1: Introduction and overview		
M	9/7	*Labor Day – no class*		
W	9/9	L2: Probabilistic modeling	1 (L1,2), I	
M	9/14	L3: Discrete random variables, information measures		
W	9/16	L4: Joint distributions, marginalization, joint information	2 (L3,4)	1
F	9/18			I
M	9/21	L5: Conditioning, Bayes' Rule, conditional information		
W	9/23	L6: Independence structure, mutual information	3 (L5,6), II	2
M	9/28	L7: Decision-making, most probable configurations, MAP rule		
W	9/30	L8: Graphical models, trees, hidden Markov models	4 (L7,8)	3
F	10/2	*Add Date*		II
M	10/5	L9: Belief propagation: sum-product, forward-backward alg's		
W	10/7	L10: Most probable configurations: max-product, Viterbi alg's	5 (L9,10), III	4
M	10/12	*Columbus Day – lecture shifted to 10/13*		
T	10/13	L11: Parameter estimation, Maximum Likelihood method *No recitation (Monday schedule)*		
W	10/14	L12: Learning tree models, Naïve Bayes Models, HMMs	6 (L11,12)	5
M	10/19	L13: Learning classifiers; logistic and softmax regression Optional Quiz Review		
W	10/21	**Midterm Quiz** (through L12), 7:30-9:30pm, *no class*		
M	10/26	L14: Markov/Chebyshev bounds, law of large numbers		
W	10/28	L15: Typical sets, compression and hashing	7 (L13-15)	6
F	10/30		IV	III
M	11/2	L16: Joint typicality, classification and communication		
W	11/4	L17: Atypical sequences, large deviations and max-entropy	8 (L16,17)	7
M	11/9	L18: Markov chains and random walks		
W	11/11	*Veterans Day — no class*		
F	11/13			IV
M	11/16	L19: Sampling and approximate inference		
W	11/18	L20: Markov chain Monte Carlo and Gibbs sampling *Drop Date*	9 (L18-20), V	8
11/21-29		*Thanksgiving Week Break – no classes*		
M	11/30	L21: Continuous random variables, information measures		
W	12/2	L22: Joint PDFs; continuous inference	10 (L21-24)	9
F	12/4			V
M	12/7	L23: Jointly Gaussian random variables, innovations		
W	12/9	L24: Gaussian inference, modeling, asymptotics **Final Exam**, *Finals Week, 12/14–12/18*		