Massachusetts Institute of Technology
Department of Electrical Engineering and Computer Science
6.438 Algorithms for Inference
Fall 2019

Information Sheet

Lecturers

Guy Bresler	Gregory W. Wornell
Office: 32-D672	Office: 36-677
Tel: x4-1549	Tel: x3-2297
E-mail: `guy@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

Brian Hie	E-mail: `brianhie@mit.edu`
Sanjeev Murty	E-mail: `smurty@mit.edu`
Basil Saeed	E-mail: `bsaeed@mit.edu`

Lectures:	Tuesday and Thursday, 9:30am–11am, Room 32-123
Recitations:	Friday, 10–11am or 11am–noon, Room 1-190
Office Hours:	Days, times and locations posted on stellar site.

Welcome to 6.438 !

This is a graduate-level introduction to the principles of statistical inference with probabilistic models defined using graphical representations. 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 networks and control. Moreover, it is the companion course to 6.437 (Inference and Information) which is offered each Spring; 6.437 and 6.438 may be taken in either order.

It is worth stressing that 6.438 is an *introductory* graduate subject: it is not an advanced graduate subject for students who already have a mastery of statistical inference algorithms, yet want to understand such material at an even more sophisticated level. That said, the structure of this subject will be somewhat different than other such introductions to the topic.

Ultimately, the subject is about teaching you contemporary approaches to—and perspectives on—problems of statistical inference and learning. The development of the material that forms the basis for this subject has historically been very much driven by applications. 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 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 class 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 prerequisites are 6.008, 6.041, 6.436, or 18.600, and 18.06, or their equivalents. Ultimately, what we require is *fluency* with both basic quantitative probabilistic analysis and linear algebra, together with some subsequent solid exposure to the engineering application of both. 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

We will provide lecture notes whose content is under active development. These notes will be the primary reading material. You will find some sections of these notes more mature than others as we develop and refine the treatment of different topics, and we will count on you to help us identify bugs!

Also, some of the notes (or sections thereof) will specifically be indicated as "optional" reading. These notes develop aspects of the material that we will not be covering this semester, but which you may enjoy reading about, depending on your interests and as your time permits. During the semester we will not assume you have read any of the optional notes or sections.

You may also find sections of the following books to be useful and more in-depth auxiliary references for parts of the term.

C. M. Bishop, *Pattern Recognition and Machine Learning*, Springer, 2006.

M. I. Jordan, *Introduction to Probabilistic Graphical Models*, Lecture notes (on class web site).

D. Koller and N. Friedman, *Probabilistic Graphical Models: Principles and Techniques,* MIT Press, 2009.

D. J. C. MacKay, *Information Theory, Inference, and Learning Algorithms,* Cambridge University Press, UK, 2003. (on class web site; also available online at
`http://www.inference.phy.cam.ac.uk/mackay/itila/book.html`)

M. Mezard and A. Montanari, *Information, Physics, and Computation*, Oxford University Press, 2009.

We suggest you initially hold off purchasing any books while you are gauging their usefulness to you. These books are available in the MIT Barker Library.

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

Finally, we kindly ask that you restrict course materials, including supplementary reading, for your personal use only, as we have specifically been asked to not share or otherwise distribute any portions of them to others outside the subject.

Problem Sets

There will be approximately 10 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.

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.

Computational Problems

Some of the problem sets will involve a computational component, to help you explore additional aspects of the material. The officially supported prototyping platform is Python, but you're welcome to use others like Matlab so long as you don't require assistance from the staff in converting any Python-based code or data that may be provided. These problems will be graded separately from the rest of the problem set that they appear on. As with the other problems on the homework, moderate collaboration with one or two classmates is permitted provided the code and writeup you produce is your own.

Exams

There will be two quizzes in the subject. The first quiz will be held Wednesday evening, October 23, 7-10pm, in Walker Memorial (50-340). The second quiz will be held during finals week on Monday, December 16, 1:30-4:30pm, in duPont Gym. 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	35%
Final Quiz	40%
Problem Sets	15%
Computational Problems	10%

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.

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 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/6.438`

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 6.438, 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 6.438 web site. Once you enroll yourself, you can fully participate in the discussions.

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

`6.438-staff@mit.edu`

Tentative Syllabus and Schedule

Homework				Lectures	
Out	**Due**	**Day**	**Date**	**#**	**Lecture Material**
1		R	9/05	1	Introduction and overview
		T	9/10	2	Directed graphical models
		R	9/12	3	Undirected graphical models and factor graphs
2	1	T	9/17	4	Marginalization and the elimination algorithm
		R	9/19	5	Computing beliefs on trees: sum-product algorithm
3	2	T	9/24	6	MAP elimination and max-product algorithm
		R	9/26	7	HMMs; forward-backward, Viterbi alg's; FG inference
4	3	T	10/1	8	Junction tree algorithm
		R	10/3	9	Loopy belief propagation I
		F	10/4		ADD DATE
5	4	T	10/8	10	Loopy belief propagation II
		R	10/10	11	Computational equivalence of inference tasks
		T	10/15		NO CLASS - Monday schedule (Columbus Day)
		R	10/17	12	Variational methods I
6	5	T	10/22	13	Variational methods II
		W	10/23	Q	Midterm Quiz, 7-10pm (through L12 and PS5)
		R	10/24		NO CLASS
		T	10/29	14	Sampling methods I
		R	10/31	15	Sampling methods II
7	6	T	11/5	16	Gaussian graphical models and HMMs
		R	11/7	17	Gaussian BP and Kalman Filtering
8	7	T	11/12		NO CLASS
		R	11/14	18	Sampling methods III: particle filtering
9	8	T	11/19	19	Parameterization, parameter estimation
		W	11/20		DROP DATE
		R	11/21	20	Structure learning
10	9	T	11/26	21	Learning Gaussian and related graphical models
		R	11/28		NO CLASS (Thanksgiving Vacation)
		T	12/3	22	Maximum entropy modeling and exponential families
		R	12/5	23	Modal decompositions
		F	12/6		LAST ASSIGNMENT DUE DATE
		T	12/10	24	Causal inference
			12/16	Q	Final Quiz, 1:30-4:30pm (through L25 and PS9)