

Information Sheet

Lecturer

Gregory W. Wornell
Office: 36-677
Tel: 253-2297
E-mail: gww@mit.edu

Administrative Assistant

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

Teaching Assistant

Christina Lee
E-mail: celee@mit.edu

Teaching Assistant

Donggeek Shin
E-mail: dshin@mit.edu

Lectures: Tuesday and Thursday, 9:30am–11am, Room 4-370
Recitations: Friday, 10-11am or 11am-noon, Room 8-119
Office hours: Days, times, and locations posted on web 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 designed as a core graduate subject for students in the relevant subfields of both Area I and Area II. The material in this course constitutes a common foundation for work in machine learning, signal processing, artificial intelligence, computer vision, control, and communication. 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 have 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. 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.

Prerequisites

The official prerequisites are 6.008 or 6.041 or 6.436, 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 evolving lecture notes whose content continues to mature with each year's offering of the subject. These notes will be the primary reading material.

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>)

We suggest you initially hold off purchasing any books while you are gauging their usefulness to you. We have placed these books on reserve at the MIT libraries (Barker).

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 9 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. Some of the problem sets will involve a MATLAB component, to help you explore additional 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.

Project

In addition to the problem sets, the class will include a two-part MATLAB-based class project. Part I will be a guided exercise, while the continuation Part II will be a more open-ended challenge. The project is intended to be educational, interesting, and fun! The project will go out on Tuesday, November 10 and Part I will be due the following week. There will be no problem set issued during Part I of the project. Part II will be due on Friday, December 4. You may want to make a note of the project dates in your schedule now, to help you with planning your time and coordinating with other classes during the semester.

The project details will be announced closer to the time it goes out, but plan for an engaging experience with the material.

Exams

There will be two quizzes in the subject. The first quiz will be held Wednesday evening, October 21, 7-10pm at a location to be announced when a room has been assigned by the registrar. The second quiz will be held during finals week, and is scheduled by the registrar. 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	40%
Final Quiz	40%
Project	5%
Homework	15%

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 problem sets and the project.

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.

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

6.438-staff@mit.edu

In addition, we will also use **piazza** 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 is great way to clarify 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 class web site.

Tentative Syllabus and Schedule

Homework		Lectures			
Out	Due	Day	Date	#	Lecture Material
1		R	9/10	1	Introduction, overview, preliminaries
2	1	T	9/15	2	Directed graphical models
		R	9/17	3	Undirected graphical models
3	2	T	9/22	4	Factor graphs; generating and converting graphs
		R	9/24	5	Minimal I-maps, chordal graphs, trees, Markov chains
4	3	T	9/29	6	Gaussian graphical models
		R	10/1	7	Computing beliefs on graphs: elimination algorithm
5	4	T	10/6	8	Computing beliefs on trees: sum-product algorithm
		R	10/8	9	Computing beliefs on HMMs: forward-backward algorithm
		F	10/9		ADD DATE
		T	10/13		NO CLASS - Monday schedule (Columbus Day)
		R	10/15	10	Sum-product on factor trees, MAP elimination
6	5	T	10/20	11	The max-product algorithm; Viterbi algorithm
		W	10/21	Q	Midterm Quiz, 7-10pm (through L10 and PS5)
		R	10/22		NO CLASS
		T	10/27	12	Gaussian Belief Propagation
		R	10/29	13	Gaussian HMMs And Kalman Filtering
7	6	T	11/3	14	The junction tree algorithm
		R	11/5	15	Loopy belief propagation - Part I
I,II	7	T	11/10	16	Loopy belief propagation - Part II
		R	11/12	17	Variational inference
8	I	T	11/17	18	MCMC methods and approximate MAP
		W	11/18		DROP DATE
		R	11/19	19	Approximate inference by particle methods
9	8	T	11/24	20	Parameter estimation—directed graphical models
		R	11/26		NO CLASS (Thanksgiving Vacation)
		T	12/1	21	Parameter estimation—undirected graphical models
		R	12/3	22	Estimating directed graphical model structure
	II	F	12/4		LAST ASSIGNMENT DUE DATE
		T	12/8	23	Estimating undirected graphical model structure
		R	12/10	24	TBD
		W	12/14-18	Q	Final Quiz (through L24 and PS9)