CSC 412/2506
Probabilistic Graphical Models

Lectures:	Tuesdays 3-5
Lecture Room:	MP 137
Instructors:	Richard Zemel
Instructor email:	<csc412prof@cs.toronto.edu>
Office hours:	Thursday 4-5 Pratt 290D
Teaching Assistants:	Kaustav Kundu, Yujia Li, Wenjie Luo
TA email:	<csc412ta@cs.toronto.edu>
Tutorials:	Thursdays 1-2
Tutorial Room:	Same as lecture
Class URL:	www.cs.toronto.edu/~zemel/Courses/CS412.html

Overview

A senior undergraduate/graduate class on graphical models and probabilistic networks in AI. Representing uncertain knowledge using probability and other formalisms. Qualitative and quantitative specification of probability distributions using graphical models. Algorithms for inference with graphical models. Statistical approaches and algorithms for learning models from experience. Examples will be given of applications of these models in various areas of artificial intelligence.

Pre-requisites

You should have taken an introductory machine learning course. You should understand basic probability and statistics, and college-level algebra and calculus. For example it is expected that you know about standard probability distributions (Gaussians, Poisson), and also how to calculate derivatives. Knowledge of linear algebra is also expected, and knowledge of mathematics underlying probability models will be useful. For the programming assignments, you should have some background in programming, and it would be helpful if you know Matlab or Python.

Readings

There is no required textbook for this course, but we strongly recommend purchasing *Machine Learning: A Probabilistic Perspective* by Kevin Murphy. We will specify readings

from this book. We will also provide pointers to other books and online resources.

Course requirements and grading

The format of the class will be lecture, with some discussion. I strongly encourage interaction and questions. There are assigned readings for each lecture that are intended to prepare you to participate in the class discussion for that day.

The grading in the class will be divided up as follows:

Assignments	50%
Mid-Term Test	20%
Project	30%

There will be three assignments; the first two are worth 15% of your grade and the last one is worth 20%.

Homework assignments

You will be asked to summarize your work, and analyze the results, in brief (3-4 page) write ups. The implementations may be done in any language, but Matlab or Python is recommended. A brief tutorial on Matlab is available from the course web-site.

Collaboration on the assignments is not allowed. Each student is responsible for his or her own work. Discussion of assignments and programs should be limited to clarification of the handout itself, and should not involve any sharing of pseudocode or code or simulation results. Violation of this policy is grounds for a semester grade of F, in accordance with university regulations.

The schedule of assignments is included in the syllabus. Assignments are due at the beginning of class/tutorial on the due date. Because they may be discussed in class that day, it is important that you have completed them by that day. Assignments handed in late but before 5 pm of that day will be penalized by 5% (i.e., total points multiplied by 0.95); a late penalty of 10% per day will be assessed thereafter. Extensions will be granted only in special situations, and you will need a Student Medical Certificate or a written request approved by the instructor at least one week before the due date.

Tests

There will be a mid-term in tutorial on February 25^{th}, which will be a closed book test on all material covered up to that point in the lectures, tutorials, required readings, and

assignments. You will not be responsible for topics in the reading not covered in any of these.

Project

The idea of the project is to give you some experience trying to do a piece of original research in machine learning and writing up your results in a paper style format. What we expect to see is an idea/task that you describe clearly, relate to existing work, implement and test on a dataset. To do this you will need to write code, run it on some data, make some figures, read a few background papers, collect some references, and write a few pages describing your task, the algorithm(s) you used and the results you obtained. As a rough rule of thumb, spend about one week's worth of work (spread out over a longer time to allow the computers to do some work in the interim!), and about a day writing it up after that. Projects can be done individually, or in pairs. We encourage you to work in pairs, but of course, the expectations will be higher for pair projects.

Attendance

We expect students to attend all classes, and all tutorials. This is especially important because we will cover material in class that is not included in the textbook. Also, the tutorials will not only be for review and answering questions, but new material will also be covered.

Electronic Communication

If you have questions about the assignments, you should send email to the TA account, and cc the instructors on it. You should include your full name in the email, and it will also be useful to include your CDF account name and/or student number. Feel free to email the instructors with questions or comments about the material covered in the course, or other related questions.

For questions about marks on the assignments, please first contact the TA. Questions about the exams should be addressed to the instructor.

CLASS SCHEDULE

Shown below are the topics for lectures, and the dates that each assignment will be handed out and is due. The notes from each lecture and tutorial will be available on the class web-site the day of the class meeting. All of these are subject to change.

Date	Topic	Assignments
1. Jan 12	Introduction to inference & learning	
2. Jan 19	Simple classifiers	A1 out
3. Jan 26	Undirected graphical models	
4. Feb 2	Directed graphical models	
5. Feb 9	Exact inference	A1 due; A2 out
Feb 16	*Reading week: No class*	
6. Feb 23	Variational inference	Project proposal due
Feb 25	*Mid-term*	
7. Mar 1	Sampling	A2 due; A3 out
8. Mar 8	State space models	
9. Mar 15	Conditional random fields	
10. Mar 22	Gaussian Processes	A3 due
11. Mar 29	Boltzmann Machines	
12. Apr 5	Project Presentations	
Apr 20		Projects due