CSC 2506 — Probabilistic Reasoning — Spring 2000

Web page: `http://www.cs.toronto.edu/~radford/csc2506/`

Instructor: Radford M. Neal, Office: SS 6016A, Phone: 978-4970, E-mail: radford@cs

Textbook:

R. G. Cowell, A. P. Dawid, S. L. Lauritzen, and D. J. Spiegelhalter, *Probabilistic Networks and Expert Systems*.

Copies of some papers will be distributed as well.

Also recommended:

J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*.

B. J. Frey, *Graphical Models for Machine Learning and Digital Communication*.

M. I. Jordan (editor), *Learning in Graphical Models*.

The proceedings of the annual conference on *Uncertainty in Artificial Intelligence* are also of interest.

Components of mark:

15%	Assignment
25%	First test
25%	Second test
35%	Project

Assignment:

This will involve some probabilistic model design, and some simple programming.

Tests:

These will cover the theoretical material, in a fairly straightforward way.

Possible projects:

- Design of a probabilistic model for some domain
- Implementation and evaluation of probabilistic inference algorithms
- Theoretical analysis of a model or algorithm
- Review of the literature on some topic
- Anything else that is relevant and reasonable

Some probabilistic modeling software will be available, which could be used for the project. The project will be due in April. Lengthy extensions will **not** be given, for your sake as well as mine.

Tentative Syllabus and Schedule

Jan	3	The nature of probability. Sketch of graphical models. Axioms of probability. Conditional probability. Independence.
	10	Random variables. Conditional independence. Directed graphical models. Noisy-OR and logistic models. Undirected graphical models. Chain graph models.
	17	Determining conditional independence. Examples of model creation. Probabilistic inference in directed and undirected tree models. Application: Hidden Markov models. *Assignment handed out*
	24	Factor graphs. Inference in polytree models. Can cycles be ignored? Application: Low-density parity-check codes.
	31	Creating junction trees: moralization, and triangulation. Probabilistic inference using junction trees.
Feb	7	Other views of inference: eliminating cycles, eliminating variables. Review. *Assignment due*
	14	READING WEEK
	21	Using graphical modeling software. *First Test*
	28	Gaussian and conditionally-Gaussian models. The Kalman filter. *Project proposals due*
Mar	6	Monte Carlo methods. Importance sampling. Gibbs sampling.
	13	Learning models from data. Maximum likelihood and Bayesian methods. Learning from complete data. Gibbs sampling for Bayesian learning.
	20	Learning form incomplete data (unobserved or latent variables). The EM algorithm for maximum likelihood. Gibbs sampling.
	27	Causal inference. *Second Test*
Apr	3	*Project presentations*