Probabilistic graphical models
(Advanced Topics in Artificial Intelligence)

CS 3710 – Fall 2005

Lecture meeting time: Monday, Wednesday**:** 4:00PM-5:20 PM
Classroom: 5313 Sennott Square (SENSQ)

Instructor:	**Milos Hauskrecht**
Office:	5329 Sennott Square Building
Office Hours:	Tuesday 3:00-4:00pm, Wednesday 1:00-2:00pm
Phone:	(412) 624–8845
Email:	milos@cs.pitt.edu
Course Web page:	http://www.cs.pitt.edu/~milos/courses/cs3710/

Course Description:

The Advanced Topics in AI is a seminar course and this year it will focus on probabilistic graphical models. Probabilistic graphical models are used to model stochasticity (uncertainty) in the world and are very popular in AI and machine learning. The course will cover two classes of graphical models: Bayesian belief networks (also called directed graphical models) and Markov Random Fields (undirected models). After introducing the two frameworks the course will focus on recent advances in inferences and learning with graphical models, including topics such as loopy belief propagation, variational approximations, conditional Markov random fields and others.

Prerequisites: graduate level AI (CS 2710 or its equivalent**),** or the permission of the instructor. Machine Learning course (CS 2750) is a plus.

Readings:

- Michael Jordan. Introduction to Graphical Models. In preparation.
- Daphne Koller and Nir Friedman. Bayesian Networks and Beyond. In preparation.
- Conference papers and journal articles

Other books related to the course:

- Judea Pearl. Probabilistic Reasoning in Intelligent Systems. Morgan Kaufman.
- Finn Jensen. An introduction to Bayesian Networks. Springer-Verlag.
- S. Lauritzen. Graphical Models. Oxford University Press.
- David J.C. Mackay. Information theory, inference, and learning algorithms. Cambridge, UK: Cambridge University Press.

Requirements:

- Midterm project (assigned – all students work on the same topic)
- Final project (individual)

Policy on Cheating

All the work in this course should be done independently. **Collaborations on *quizzes, exams* and *homework assignments* are not permitted.** Cheating and any other anti-intellectual behavior, including giving your work to someone else, will be dealt with severely. If you feel you may have violated the rules speak to us as soon as possible.

Please make sure you read, understand and abide by the Academic Integrity Code for the Faculty and College of Arts and Sciences (http://www.fcas.pitt.edu/academicintegrity.html).

Students With Disabilities

If you have a disability for which you are or may be requesting an accommodation, you are encouraged to contact both your instructor and Disability Resources and Services, 216 William Pitt Union, (412) 648-7890/(412) 383-7355 (TTY), as early as possible in the term. DRS will verify your disability and determine reasonable accommodations for this course.

Tentative syllabus:

- **Bayesian belief networks:**
  - Representation
  - Independence and conditional independence
  - Partial independence and other structure
- **Exact inference in BBN:**
  - Variable elimination
  - Pearl's algorithm
  - Junction tree
  - Recursive decomposition
  - Using additional structure
- **Approximate inference:**
  - Monte Carlo approximations
  - Loopy belief propagation
  - Variational methods
- **Learning of BBNs:**
  - learning parameters
  - learning structure
  - Bayesian averaging
  - EM (learning with hidden variables and missing values)
  - structural EM
- **Dynamic belief networks**
  - Particle filtering

- **Markov random fields (Markov networks)**
  - Representation (potentials)
  - Independence and conditional independence
  - Trees
  - Boltzman machines
  - Conditional Markov random fields
- **Inference in Markov networks**
- **Learning Markov networks:**
  - Iterative proportional fitting
  - Cluster variational methods
  - Other approximations
- **Relational graphical models**
- ...