CS573: Probabilistic Reasoning

The chief objective is to teach modern methods of probabilistic reasoning that are commonly used in many parts of computer science, including but not limited to artificial intelligence. Such methods have become extremely important and transforming the approach to a great variety of computational problems, in the field of computer science itself, and broadly across many application fields.

The course is aimed at first or second year M.S. and Ph.D. students in Computer Science (CS) though students in related disciplines such as Electrical Engineering (EE) and Industrial and Systems Engineering (ISE) may also find it beneficial.

The objective of the course is to provide an in-depth understanding of the topics covered. After the course, students should be able to read the relevant research literature as well as apply the methods to the myriad applications in computer science and other disciplines where reasoning (inferencing and decision making) under uncertainty is important.

Recommended preparation: undergraduate training in probability theory (such as in Math 208x); basic knowledge of computer programming, data structures and algorithms (such as found in CS 201 and 301).

Grading: Based on exams and assignments. There will be a total of five or six assignments. Most of the problems will be “quantitative” or “mathematical”, requiring students to apply the learned problems to specific examples or new results. The assignments will also contain programming exercises that involve implementation and testing of algorithms that are too complex to apply manually. The combined weights of the assignments will be 30% of the course total; two in-class, closed book exams will count for 30% of the grade each. Remaining 10% will be based on class attendance and participation. Neither attendance nor assignments are optional.

Instructor: Prof. Nevatia.

Schedule: MW 11-12:20

Textbook:

Probabilistic Graphical Models, by Daphne Koller and Nir Friedman, MIT Press, 2009.

Class Schedule:

Following is an expected schedule of topics and the corresponding pages in the textbook. This estimate is based on a previous offering of the course. Topics in this semester may change to some extent, based on the pace of the class and the interest of the students.

Class #	Topic	Textbook pages
1	Class policies, Course Objectives, Introduction	--
2	Overview of Probability Theory	15-36
3	Bayes Networks	45-67
4	Independence, I-Maps	68-92
5	Undirected Graphical Models	103-128
6	Bayes Networks and Markov Networks	128-152
7	Local Models	157-178
8	Template Based Representations	199-216, 228-232
9	Exact Inference: Variable Elimination;	287-306
10	Clique Trees, Belief Propagation	345-369
11	Tree Construction, Intro to Optimization	369-374, 381-399
12	Approximate Inference: Sampling	487-505
13	Markov Chains	505-526
14	MAP Inference	551-572
15	Exam 1 (tentative)	--
16	Inference in Temporal Models	651-675
17	Learning Graphical Models : Intro	697-714
18	Parameter Estimation	717-741
19	Bayesian Networks and Shared Parameters	742-763
20	Structure Learning and search	783-837
21	Partially Observed Data, Gradient Descent	850-868
22	EM, Hidden Variables	868-893, 925-930
23	Undirected Models	943-961
24	Causality	1009-1031
25	Utility Functions	1057-1069
26	Decision Problems	1069-1079, 1083-1096
27	Expected Utility, Value of Information	1098-1124
28	Exam 2	--