EECS 395/495: Special Topics in Computer Science: Probabilistic Graphical Models Fall 2011

Introductions

- Professor: Doug Downey
- Course web site:
 - www.cs.northwestern.edu/~ddowney/courses/395/
 - (linked off prof. home page)

Introductions

- Your Name
- Major/Degree
- One thing you hope to learn
- Do you want "more math" or "less math"?

Logistics

- Grading
 - Homework (50%)
 - Handed out in weeks 1, 2, 4, 5
 - Exercises and programming
 - Midterm (25%)
 - Week 7
 - A lot like the homework
 - Paper Presentation (15%)
 - Class Participation (10%)

Textbook

 D. Koller & N. Friedman, *Probabilistic Graphical Models: Principles and Techniques* MIT Press, 2009.



What's going on

- Artificial Intelligence
 - tremendous success in domains without a lot of uncertainty (e.g. chess)
 - But in the real world, uncertainty reigns
- We are awash in data

A crisis and an opportunity

 How can we deal with uncertainty? And how can we exploit massive bodies of data?

What the course is about

- Probabilistic Models
 - Deal with uncertainty: assign degree of confidence that different events will occur
- Probabilistic Graphical Models
 - Graph-based representation
 - Compactly encode complexity



Goals

- Learn how to:
 - Build probabilistic models from data
 - Use the models to perform tasks
 - Recognize opportunities for using models

Compared with Other Courses (1 of 2)

- EECS 349 vs. this class
 - EECS 349 is a prerequisite for this class
 - EECS 349 focuses on learning *functions*, we learn *distributions*
 - A single distribution captures many different functions!
 - EECS 349 more algorithmic, this class more mathematical

Compared with Other Courses (2 of 2)

- Statistics vs. this class
 - A few variables vs. tens of thousands
 - Continuous vs. discrete variables
 - Our focus: computational issues and applications
 - How can we scale to huge, multivariate data sets?
 - When and where are graphical models useful?

Applications

- Almost anything!
- E.g.,
 - Computational Biology
 - Robotics
 - Vision
 - Human-Computer Interaction
 - Networks and Systems
 - Information Retrieval/Web Search
 - Etc., etc.

Topics

- Basics of Probability and Statistical Estimation (briefly)
- Representing Probability Distributions as Graphs
 - Directed ("Bayes Nets") and Undirected ("Markov Nets")
- Working with Probabilistic Graphical Models
 - Inference: answering queries with a model
 - Learning: acquiring models from data
- Sequential Models (Hidden Markov Models), Statistical Language Models, Active Learning, ...

Paper Presentations

- Teams of 1-2
- Cover:
 - Key ideas
 - Questions for discussion
 - (might take less than 80 minutes)
- I'll do the first one
- Paper list to appear this week