EECS 395/495-21:
Special Topics in Computer Science

Winter 2010

Introductions

- Professor: Doug Downey
- TA: Jiang Xu
- Course web site:
  - www.cs.northwestern.edu/~ddowney/courses/395_Winter01/
  - (up soon – will be linked off prof. home page)
- Watch for e-mail (to addr. in Blackboard)

Introductions

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

Logistics

- Grading
  - Homework (75%)
    - Handed out in weeks 1, 2, 4, 6, and 8
    - Exercises and programming
    - Several hands-on exercises will utilize a “target domain”
      - Pick a task of your choice
      - Hmwk exercises: acquire data, build models, analyze results
      - You can partner arbitrarily
  - Final (25%)
    - A lot like the homework

Textbook

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

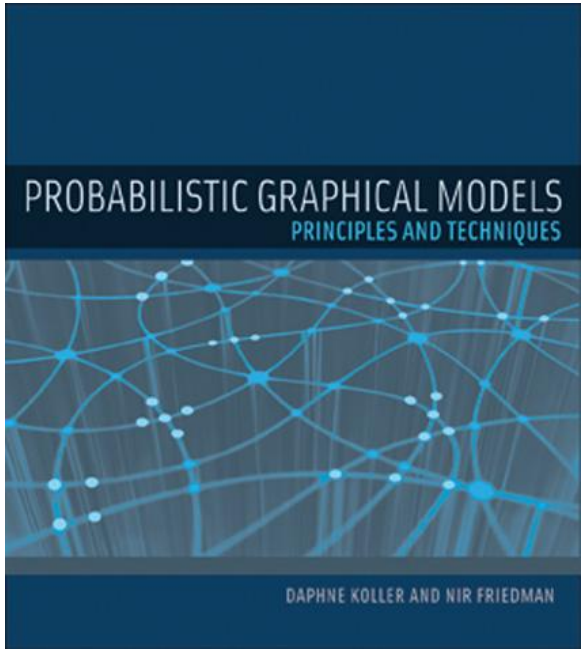

What's going on

- 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

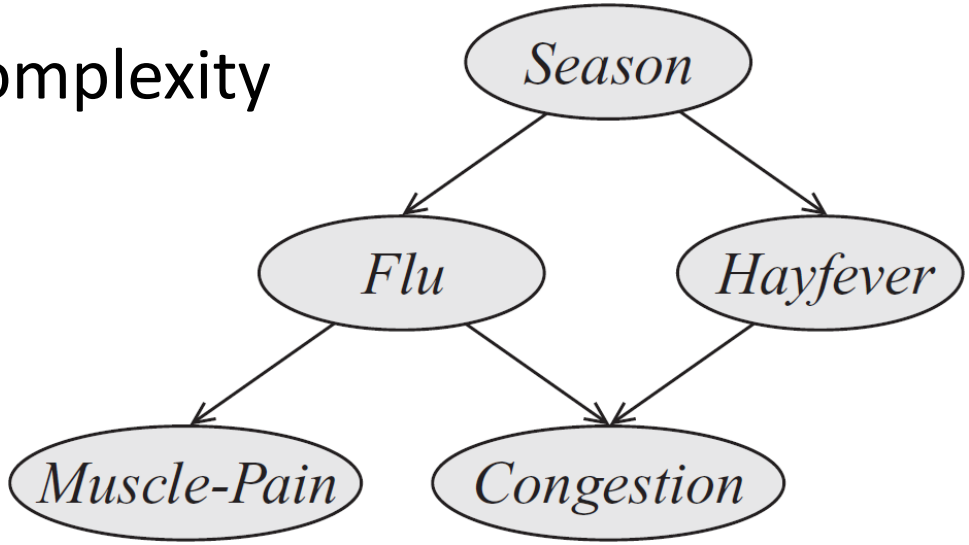

Goals

- Learn how to:
  - Build probabilistic models from data
  - Use the models to do work
  - Recognize opportunities for using models
- My focus
  - Interaction
  - Evaluation

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)
- Time Allowing:
  - Active Learning, Decision Theory, Relational Models