EECS 598-04 Special Topic

Probabilistic Graphical Models for Vision and Beyond

Winter 2015: MW 9-10:30
EECS 1003

Course Overview: Probabilistic graphical models have emerged as a powerful formalism for leveraging principled probability theory and discrete structured data representations to model large-scale problems involving hundreds or even thousands of often high-dimensional variables. Computer vision provides a set of such large-scale, concrete problems well-suited for probabilistic graphical models. The course will cover probabilistic graphical models in detail starting from the basics and pushing through contemporary results. There will be an emphasis on driving problem formulations from computer vision but our coverage will be broad; connections to other application areas will be discussed when plausible. No prior course in computer vision is needed (although it will help in familiarity with some terminology).

Syllabus: Foundations of graphical modeling will begin the lectures followed by a thorough discussion of causal (Bayesian) and acausal (Markov) graphical models. Thereotical bases of these two paradigms will be covered (e.g., Gibbs distributions, Hammersley-Clifford theorem). Learning and inference methods (exact, approximate, discrete and stochastic methods) will be discussed in detail. Practical considerations for all three course components (modeling, learning and inference) will be covered both through course discussion and assignments.

Course Structure: The course will be meet twice per week. Each week one lecture will be presented on a specific topic and one interactive discussion will be conducted based on that topic and assigned course work. 50% of the grade will be based on this assigned course work (essentially these are weekly or biweekly problem sets). The other 50% of the grade will be based on a course project. Course projects may be applied to any application area with instructor approval.

Prerequisites: EECS 501 or graduate-level proficiency with probability and statistics. In addition, at least one of these is necessary (but the more you have, the better):
- EECS 445/545 or a prior course in machine learning.
- EECS 280 or proficiency in programming with data structures.
- EECS 442/542 or a prior course in computer vision.

Permission of instructor can be sought to waive either of these.

Instructor: Prof. Jason Corso, Electrical Engineering and Computer Science

Website: http://web.eecs.umich.edu/~jjcorso/t/598W15/