EECS 598-04 Probabilistic Graphical Models

Electrical Engineering and Computer Science

University of Michigan

Syllabus for Winter 2015

Last updated: 6 January 2015

Instructor: Jason Corso (jjcorso)

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

Syllabus: http://web.eecs.umich.edu/~jjcorso/t/598W15/files/598W15_syllabus.pdf.

Meeting Times: MW 0900–1030

Location: 1003 EECS

Office Hours: Monday 1500-1600 and Wednesday 1030-1130 (4227 EECS)

Course Information Flow and A Note On Contacting The Instructor: This courses uses CTools, Piazza and the instructor's website.

- **CTools** will be used for some minor things such as distributing protected material, accessing grades. CTools Link is https://ctools.umich.edu/portal/site/b9118a40-6399-473a-9321-6a89e6a131b9.
- **Instructor's website** (http://web.eecs.umich.edu/~jjcorso/t/598W15) will hold the schedule.
- **Piazza** is used for announcements, news and discussions. The piazza course website is http://piazza.com/umich/winter2015/eecs59804/home (this link is also in the CTools site). Students should ensure they are enrolled in the course. Nearly all questions you have about the course, both logistical and technical should be posted to piazza (after you have already checked piazza to ensure the same question has not already been answered). Only in the event of a concern of privacy, should you directly email the instructor.

Main Course Material

Course Description: 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 inter-related variables. The course will cover probabilistic graphical models in detail starting from the basics and pushing through contemporary results. 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 (representation, learning and inference) will be covered both through course discussion and assignments.

There will be an emphasis on driving problem formulations from computer vision since it provides ample problems well-suited for probabilistic graphical models but our coverage will be broad and include robotics, natural language processing, computation finance and other areas. No prior course in computer vision or other application area is needed (although it will help in familiarity with some terminology).

Prerequisites: EECS 501 or graduate-level proficiency with probability and statistics. EECS 445/545 or a prior course in machine learning. Computer programming proficiency (data structures and algorithms course such as EECS 280/281; Matlab and C++).

Course Goals: After taking the course, the student should

1. have a clear understanding of the foundations of probabilistic graphical models as a framework for representation, learning and inference in complex systems over many variables with high degrees of uncertainty,
2. have gained practical experience of probabilistic graphical models, and
3. be set-up to use probabilistic graphical modeling in their own research and/or have a strong enough foundation to pursue research on probabilistic graphical models themselves.

These goals are evaluated through the in-class discussion, presentation, homeworks and a project.

Textbooks: The required textbook for the course is

- Koller and Friedman *Probabilistic Graphical Models: Principles and Techniques*, MIT Press 2009

The textbook has a website: pgm.stanford.edu

Programming and Software: This course strives to find a balance between theoretical and practical considerations of probabilistic graphical models. To make the practical side of the course plausible across a larger class, we will collectively use a common graphical modeling library:

- OpenGM is a C++ template library for discrete factor graph models and distributive operations on these models. http://hci.iwr.uni-heidelberg.de/opengm2/

Portions of OpenGM will be used out of the box, portions will be reimplemented and various new contributions may be made. OpenGM is selected over other options, including those natively in higher level languages like Matlab, because of its breadth. Although no full Matlab interface is provided, one can generate OpenGM readable files in Matlab and vice-versa; so for various projects, one can use Matlab and OpenGM together.

Course Work and Evaluation

Homeworks (25%) There will be two analytical homework assignments given in the first half of the semester. They will be mathematical problems to test whether students are understanding the foundational topics. These assignments will be done independently by each student.

In-Class Presentation (25%) Each student will give one-to-three in-class presentations in class. The presentation topic will be specified by the instructor and vary from lecturing on a particular model or algorithm, discussing a particular implementation of a model or algorithm, or presenting a paper that applies ideas in probabilistic graphical models to various domains. Depending on the size of the course, these in-class presentations may be in small groups.

Project (50%) Small groups of students will select, formulate and implement a project involving probabilistic graphical models in a real-world situation, such as image understanding, robot navigation, stock-market behavior, acoustic speech recognition, etc. Topics will be selected and proposed by the group and then approved / rejected-for-revisions by the instructor. Full-scale project developments will be made through the term culminating in a conference-length paper and a public poster/demo session at the end of the term. In past semesters, such projects have gone on to be published at first-tier conference and journal venues.

More information about the project will be distributed in the second week of class.

No exams will be given in the course.

Late Work Policy: **No** late work will be accepted. Ample time will be given to complete the homeworks and project; use it wisely. This is a firm policy. Do not expect special treatment.

Regrading: Any questions about the grading of a piece of work must be raised within one week of the date that the work was returned by the GSI or the instructor. In other words, if you do not pick up your work in a timely fashion, you may forfeit your right to question the grading of your work.

Additional Information

Similar Courses at Other Institutions: (incomplete and in no important order)

- Koller (Coursera): https://www.coursera.org/course/pgm
- Krause (CalTech): http://courses.cms.caltech.edu/cs155/
- Sontag (NYU): http://cs.nyu.edu/~dsontag/courses/pgm13/
- Sudderth (Brown): http://cs.brown.edu/courses/csci2950-p/
- Lee (UW): https://courses.cs.washington.edu/courses/cse515/11sp

General Notes

If you don't understand something covered in class, ask about it right away. The only silly question is the one which is not asked. If you get a poor mark on an assignment or exam, find out why right away. Don't wait a month before asking. The instructor and GSI(s) (if associated). Don't be afraid to ask questions, or to approach the instructor or TA in class, during office hours, through the discussion board or through e-mail. This course is intended to be hard work, but it is also intended to be interesting and fun. We think computer vision is interesting and exciting, and we want to convince you of this.

Disabilities

If you think you need an accommodation for a disability, please let me know at your earliest convenience. Some aspects of this course, the assignments, the in-class activities, and the way the course is usually taught may be modified to facilitate your participation and progress. As soon as you make me aware of your needs, we can work with the Office of Services for Students with Disabilities (SSD) to help us determine appropriate academic accommodations. SSD (734-763-3000; http://ssd.umich.edu) typically recommends accommodations through a Verified Individualized Services and Accommodations (VISA) form. Any information you provide is private and confidential and will be treated as such.

Counseling Center

Your attention is called to the Counseling and Psychological Services (734-764-8312), 3100 Michigan Union. The Counseling Center staff are trained to help you deal with a wide range of issues, such as how to deal with exam-related stress and other academic and non-academic issues. Services are free and confidential and do not impact student records. Shivaun Nafsu is the CAPS consultant directly within COE: snafsu@umich.edu or 734-763-8211. Their web site is http://caps.umich.edu/.

Standards of Conduct – Behavioral Expectations

The following are classroom "etiquette" expectations:

- Attending classes and paying attention. Do not ask an instructor in class to go over material you missed by skipping a class or not concentrating.

- Not coming to class late or leaving early. If you must enter a class after lecture has clearly begun, do so quietly and do not disrupt the class by walking between the class and the instructor. Do not leave class unless it is an absolute necessity.
- Not talking with other classmates while the instructor or another student is speaking. If you have a question or a comment, please raise your hand, rather than starting a conversation about it with your neighbor.
- Showing respect and concern for others by not monopolizing class discussion. Allow others time to give their input and ask questions. Do not stray from the topic of class discussion.
- Not eating during class time.
- Turning off the bothersome electronics: cell phones, pagers, and beeper watches.
- Avoiding audible and visible signs of restlessness. These are both rude and disruptive to the rest of the class.
- Focusing on class material during class time. Sleeping, talking to others, doing work for another class, reading the newspaper, checking email, and exploring the internet are unacceptable and can be disruptive.
- Not packing bookbags or backpacks to leave until the instructor has dismissed class.

College of Engineering Honor Code

The full Engineering Honor Code is available at `http://www.eecs.umich.edu/acal/honor.html`. All students are expected to read, understand and follow the honor code.

The Honor Code outlines certain standards for ethical conduct for persons associated with the College of Engineering at the University of Michigan. The policies of the Honor Code apply to graduate and undergraduate students, faculty members, and administrators.

In 1915, the students of the College of Engineering petitioned for the establishment of an Honor Code. The Code was promptly adoped with faculty approval and has been basic to life in the College of Engineering.

The Honor Code rests upon the following principles:

- Engineers must possess personal integrity both as students and as professionals. They must be honorable people to ensure safety, health, fairness, and the proper use of available resources in their undertakings.
- Students in the College of Engineering are honorable and trustworthy persons.
- The students, faculty members, and administrators of the College of Engineering trust each other to uphold the principles of the Honor Code. They are jointly responsible for precautions against violations of its policies.
- It is dishonorable for students to receive credit for work which is not the result of their own efforts.

The Engineering Honor Code is based on the principle that students will follow all guidelines for study and prepared work set forth by the instructor, and that students can be trusted to take examinations without cheating.

Students are responsible for reporting infractions of the honor code.