GEORGIA INSTITUTE OF TECHNOLOGY

ECE/CS/ISYE 8803
Probabilistic Graphical Models in Machine Learning
Spring Semester 2018

Instructor:

Faramarz Fekri
email: fekri@ece.gatech.edu

Teaching Assistant:

TBA

Class Hours:

TR: 4:30-5:45 PM

Office Hours:

TR: After Class
10:00AM-12:00PM Fridays, (Location: My office at “Centergy One Building”, room 5238)

Course Objectives:

The course will provide students with an introduction to the theory and practice of graphical models, one of the most dominant frameworks in machine learning and artificial intelligence. The class will cover three main aspects: The core **representation**: including Bayesian and Markov networks, and dynamic Bayesian networks; **probabilistic inference algorithms**: both exact and approximate; and **learning methods**: for both the parameters and the structure of graphical models. Specifically, students will:

1. Become familiar with the most commonly used graphical model representation methods, learning and inference algorithms.
2. Gain exposure to the application of graphical models to real world problems.
3. Learn as to how one can formulate a wide range of problems with very large number of variables using the unified language of graphical models.

Textbook:

- Most of the materials are covered in the book “Probabilistic Graphical Models: Principles and Techniques,” Daphne Koller & Nir Friedman.
- Lecture Slides will be provided.
- A few topics will be based on the book “An Introduction to Probabilistic Graphical Models” in preparation by Michael I. Jordan. PDF of a few Chapters of the book (as the "duplicate notes") will be provided as the course progresses.
- Additional Reference: “Machine Learning A Probabilistic Perspective,’’ Kevin P. Murphy, MIT Press.

Course Prerequisites:

Familiarity with basic linear algebra and basic probability theory is assumed.

Grading Formula:

Project	30%
Homework	40%
Final Exam	30%

Homework:

The primary way to learn any subject is to WORK HOMEWORK PROBLEMS: as many as possible, and work them CAREFULLY. There will be approximately 4 to 5 Homeworks. They will be due at the beginning of the class on due dates (which are at least two weeks from the date it was assigned). Late homework will NOT be accepted for grading**.** Homework is to be written up and submitted individually. Working with colleagues is encouraged but simply copying someone else's solution is not acceptable and will be treated as such. Homework will be graded and solutions will be available.

Project:

The project gives you the opportunity to study in greater depth certain concepts of the course. The topic has to be linked with algorithms, concepts or methods presented in class, but beyond this requirement, the choice is quite open. In particular, it may be tailored to your interests. Projects should be done in teams of three students.

The standard class projects may contain one or all the following 3 components:

I. Article reviews around a given topic (not studied in class). This means to read and understand a specific research topic.
II. An implementation of the method studied in item I in above.
III. An experimentation with real data for the method in item II in above. This means to apply the method on real data and report your findings and observations.

The project will have 3 deliverables:

1. **Proposal and teaming:** (2%), 1 page. (Due: Feb. 15)
2. **Final Report:** (13%) At least 15 pages (single column, single space, 11 size font), (Due: Last week of the semester)
3. **15 minute Presentation (5 minutes per team member) & no less than 15 Slides:** (15%)(Due: last two weeks of the semester)

A list of potential projects will be provided to you by January 23. Ideally, you will want to pick a problem in a domain of your interest, e.g., natural language parsing, DNA sequence analysis, text information retrieval, network mining, social computing, reinforcement learning, sensor networks, etc., and formulate your problem using graphical models. You can then, for example, adapt and tailor standard inference/learning algorithms to your problem, and do a thorough performance analysis.

Attendance:

Regular attendance in class is mandatory**.**

Honor Code:

Academic dishonesty will not be tolerated. This includes cheating, lying about course matters, plagiarism, or helping others commit a violation of the Honor Code. Plagiarism includes reproducing the words of others without both the use of quotation marks and citation. Students are reminded of the obligations and expectations associated with the Georgia Tech Academic Honor Code and Student Code of Conduct, available online at www.honor.gatech.edu.

Class Web Page:

http://fekri.ece.gatech.edu/course_ece8803.html