Course Syllabus: CPCSI 691GM
Graphical Models – Spring 2013

Instructor: Benjamin M. Marlin
Email: 691gm@cs.umass.edu
Office: CS234
Office Hours: TBA

First Class: January 22, 2013
Time: TuTh 2:30-3:45
Location: CS 142

Course Description: Probabilistic graphical models are an intuitive visual language for describing the structure of joint probability distributions using graphs. They enable the compact representation and manipulation of exponentially large probability distributions, which allows them to efficiently manage the uncertainty and partial observability that commonly occur in real-world problems. As a result, graphical models have become invaluable tools in a wide range of areas from computer vision and sensor networks to natural language processing and computational biology. The aim of this course is to develop the knowledge and skills necessary to effectively design, implement and apply these models to solve real problems.

The course will cover (a) Bayesian and Markov networks and their dynamic and relational extensions; (b) exact and approximate inference methods; (c) estimation of both the parameters and structure of graphical models. Although the course is listed as a seminar, it will be taught as a regular lecture course with programming assignments and exams. Students entering the class should have good programming skills and knowledge of algorithms. Undergraduate-level knowledge of probability and statistics is recommended. 3 credits.

Textbook: The course textbook is *Probabilistic Graphical Models* (http://pgm.stanford.edu/) by Koller and Freidmann. It is strongly recommended.

Course Website: The course website is hosted on UMass's Moodle course management portal https://moodle.umass.edu/. You log into the portal using your OIT user name and password. The course website will host lecture notes, assignments, pointers to readings and demos, announcements, and discussion forums.

Announcements: Official announcements for the course will go out through the Moodle portal as email and will be automatically logged as news items. Email will be sent to your *official UMass email address* (@student.umass.edu, @cns.umass.edu, etc...). If you want to redirect this email elsewhere, login to the UMail Post Office (https://postoffice.oit.umass.edu/umpo/exec/login) and setup mail forwarding to an alternate address.

Grading Plan: The coursework will consist of homework assignments, quizzes, reading responses, in-class discussions and a final exam. The assignments will contain some written questions, but will mostly consist of programming and experimentation. Assignment solutions will be written-up as short reports. A one-paragraph response will be required for the readings. The grading scheme is below:

Homework Assignments	50%
Final Exam	25%
Quizzes	10%
Reading Responses	10%
Course Participation	5%

Course Policies:

- **Homework Submission:** Homework assignments will generally consist of a written component consisting of derivations, short answer questions and analyses as well as a programming component. The written component must be handed in at the start of class on the date it is due. Code for the programming component must be uploaded to Moodle before class on the date it is due. Code must be submitted as a single zip file (not rar, gz, tgz, bz2, etc.). In the event that you can not attend class on the date an assignment is due, you may submit a scanned or electronic copy of the written portion of your assignment by email before the start of class.
- **Late Homework:** To allow some flexibility to complete assignments while working on research or other coursework, you have a total of five free late days that you can use however you choose. You will be charged one late day for handing in an assignment within 24 hours after it is due, two late days for handing in an assignment within 48 hours after it is due, etc. Your assignment is considered late if either the written or code portions are submitted late. The late homework clock stops when both the written and code portions are submitted. The late written portion of your assignment must be handed in to the instructor in person or by email. The code portion must be uploaded to Moodle. After you have used up your late days, late homework will not count for credit except in special circumstances (ie: illness documented by a doctors note). Any unused late days will be credited as bonus points on your final grade at a conversion rate of 0.5% per late day.
- **Homework Collaboration:** You are encouraged to discuss assignments and course material with other students in person or on the course forums. However, you must show that you fully understand the solution to any homework problem arising from such collaboration by writing your own code, running your own experiments, and producing your own write-up as appropriate for the problem. You must also indicate which other students you collaborated with in formulating a solution to any homework problem.
- **External Resources:** For programing assignments, you are allowed and encouraged to use libraries and existing code at *one level of abstraction lower than the material we are covering in this course*. For example, if you are experimenting with inference in Bayesian networks, you can use an existing code library implementing data structures for graphs, but not an existing library implementing Bayesian networks. If you are experimenting with learning multivariate Gaussian models, you can use an existing code library implementing matrix operations, but not a library implementing Gaussian models. If you are in doubt about whether using a particular code library would be permitted for an assignment, ask. You must provide citations for all code and libraries that your programming assignment solutions build on. In formulating solutions to written homework problems, you may use background material other than the textbook and the lecture notes, but you must provide citations. Copying solutions from other sources is obviously not permitted.
- **Re-grading Policy:** Errors in grading of assignments and exams can occur despite the best efforts of the course staff. If you believe you've found a grading error, staple a short written request to your original assignment explaining where you believe the mistake was made and submit it to the instructor. Re-grade requests must be submitted no later than one week after the assignment is returned. Note that re-grading may result in your original grade increasing or decreasing as appropriate.
- **Course Participation:** Students are expected to attend each class and participate in discussions, which will count toward your classroom participation grade.

Important Dates – Spring 2013:

Jan 22, 2013	First day of classes
Feb 4, 2013	Add/Drop date
Feb 18, 2013	Holiday – President's Day
Feb 19, 2013	Monday class schedule
Mar 4, 2013	Last day for grad students to drop with DR
Mar 7, 2013	Last day for undergrads to drop with W
Mar 16, 2013	Spring Break Starts
Mar 25, 2013	Classes Resume
Apr 15, 2013	Holiday – Patriot's Day
Apr 17, 2013	Monday Schedule
May 1, 2013	Last Day of classes

Approximate Schedule: (Subject to change over the semester)

Lecture	Date	Topics
Lecture 1	Tu Jan 22, 2013	Introduction: Course logistics. Probability theory.
Lecture 2	Th Jan 24, 2013	Probability Calculus and Independence
Lecture 3	Tu Jan 29, 2013	Bayesian Network Representation
Lecture 4	Th Jan 31, 2013	Bayesian Network Properties
Lecture 5	Tu Feb 5, 2013	KL Divergence, Model Estimation and Learning
Lecture 6	Th Feb 7, 2013	Bayesian Network Learning
Lecture 7	Tu Feb 12, 2013	Markov Network Representation
Lecture 8	Th Feb 14, 2013	Markov Network Operations
	Tu Feb 19, 2013	No class – Monday schedule
Lecture 9	Th Feb 21, 2013	Comparing Markov and Bayesian Networks
Lecture 10	Tu Feb 26, 2013	Inference and the Variable Elimination Algorithm
Lecture 11	Th Feb 28, 2013	Variable Elimination and Factor Products and Clique Trees
Lecture 12	Tu Mar 5, 2013	Sum Product Message Passing Algorithm
Lecture 13	Th Mar 7, 2013	Sum Product Algorithm Analysis and Implementation
Lecture 14	Tu Mar 12, 2013	The Belief Propagation Algorithm
Lecture 15	Th Mar 14, 2013	Loopy Belief Propagation and Alternatives
	Tu Mar 19, 2013	No class – Spring Break
	Th Mar 21, 2013	No class – Spring break
Lecture 16	Tu Mar 26, 2013	Particle Representations and Monte Carlo Integration
Lecture 17	Th Mar 28, 2013	Introduction to Markov Chains
Lecture 18	Tu Apr 2, 2013	Markov Chain Monte Carlo Methods
Lecture 19	Th Apr 4, 2013	Markov Chain Monte Carlo Implementation
Lecture 20	Tu Apr 9, 2013	Markov Chain Monte Carlo and Learning
Lecture 21	Th Apr 11, 2013	Variational Inference
Lecture 22	Tu Apr 16, 2013	Variational Learning and the EM Algorithm
Lecture 23	Th Apr 18, 2013	Bayesian Inference
Lecture 24	Tu Apr 23, 2013	Bayesian Computation
Lecture 25	Th Apr 25, 2013	Bayesian Modeling and Conjugate Priors
Lecture 26	Tu Apr 30, 2013	Review for Final