

CSE 592: Advanced Topics in Computer Science
Graphical Models
(Fall 2009)
Department of Computer Science
Stony Brook University

Syllabus

Instructor: Luis Ortiz

Office: 2313-C Computer Science Building

Phone: (631) 632-1805

Email: leortiz@cs.sunysb.edu

Web Page: <http://www.cs.sunysb.edu/~leortiz>

Office Hours: 1:00pm-4:00pm, Wednesdays, or by appointment

General Information (2008-2010 Graduate Bulletin)

Course Description: An advanced lecture course on a new topic in computer science. The course is primarily designed for M.S. students, but can be taken by Ph.D. students as well.

Prerequisites: Limited to CSE graduate students; others, permission of instructor.

Credits: 3

Grading: ABCDF

Fulfilling Program Requirements: May be repeated for credit as the topic changes, but cannot be used more than twice to satisfy CSE major requirements for M.S.

Overview

Graphical models have revolutionized how we model and study large complex systems. The impact of graphical models has truly been broad. One can now find uses or instances of graphical models in statistics, computer science, applied mathematics, artificial intelligence, machine learning, electrical engineering, finance, psychology, cognitive science, game theory, economics, and the list keeps growing. Graphical models are used to address many problems in computer vision, robotics, (autonomous) control, reasoning and decision-making under uncertainty, coding and information theory, speech recognition, natural language processing, computational linguistics, machine translation, banking, computational biology and computer music, for example. Specific applications that rely on graphical models technology are by now wide-ranging and include expert systems in medical diagnosis, bioinformatics and system biology, forensics and genetic analysis, image processing and analysis, voice assisted/activated systems, automatic speech-to-text annotation/transcription and language translation systems, music-accompaniment systems, information retrieval, computer system troubleshooting, communications, agriculture, risk analysis, loan/credit rating, fraud detection, environmental conservation, water management and transportation, to name a few.

At the core of the representation is a graph (or network) which facilitates the modeling and interpretation of complex interactions between many entities in the system. The nature of the interactions depend on the system under study; as examples, the interactions may involve constraining, probabilistic, strategic and/or preferential relationships. The generality of the framework results

from an available toolkit of basic representations and algorithms from which one can draw to develop a specific application.

This course will cover the fundamental ideas behind graphical models in the context of various representations and the computational methods that support them.

Purpose: To introduce students to fundamental representational and algorithmic ideas behind graphical models, to expose them to state-of-the-art methods and applications in a variety of domains, and prepare them for continuing research and development work in the area

Objectives: To provide an introduction to the general area of graphical models (and their applications), as well as the typical problems addressed, the basic representations and algorithms often used, and the fundamental ideas behind them

Goals: At the end of the course, students should be able to

- describe and explain
 1. the fundamentals of graphical models and algorithms,
 2. the typical problems addressed,
 3. the different solution concepts used,
 4. the basic techniques available, and
 5. the basic tools for algorithmic analysis;
- recognize, identify and match the appropriate method for a particular model;
- summarize and present previous work and current methods;
- apply existing models and algorithms; and
- apply the common algorithmic analysis to extensions and variations.

Content

Organization: The course format involves formal lectures, discussions and presentations, some led by the students themselves. Reading material from a variety of sources, including textbooks, tutorials, classical and recent research papers, and other possible web content, will be handed out, posted or assigned to supplement the lectures.

Tentative List of Topics

Basic Topics

- **Constraint Graphical Models**

- *Framework*: constraint satisfaction problems (CSPs)
- *Model*: constraint network
- *Concepts*: primal, hidden and dual representations; compact representation and handling special constraints (e.g., binary, context-sensitivity); chordal graphs and triangulation; treewidth; (optimal) elimination orders; computational complexity
- *Problems*: finding consistent/satisfying assignments
- *Methods*: backtracking search; constraint propagation; dynamic programming (DP); variable elimination (VE); join/clique trees (JT); cutset conditioning (CC); simulated annealing and stochastic (local) search
- *Advanced Topics*: hypergraph representations; factor graphs; k -SAT; survey propagation (SP)

- **Probabilistic Graphical Models**

- *Framework*: probability theory; (Bayesian) statistics
- *Models*: Markov random fields; Markov networks; Bayesian networks
- *Special Models*: Ising models; Naive-Bayes; Gaussian networks; hybrid (mixed discrete-continuous) models; mixture models (e.g., mixture of Gaussians); QMR-DT
- *Concepts*: Gibbs distributions; conditional independence; Hammersley-Clifford Theorem; context-sensitive independence; reasoning under uncertainty; inference formulation of CSPs
- *Problems*: most-likely assignment and MAP estimation; posterior distributions and belief inference
- *Methods*: clique/junction trees; max-/sum-product algorithm; (generalized) belief propagation (BP); mean-field approximation and variational methods; stochastic simulation, likelihood weighting and (adaptive) importance sampling; Markov chain Monte Carlo (MCMC), Metropolis-Hastings and the Gibbs sampler
- *Learning from data*: complete vs. missing data; learning model parameters; learning model structure; model selection; maximum-likelihood (ML) and Bayesian statistical estimation; expectation-maximization (EM); gradient descent and local search
- *Advanced topics*: linear programming (LP) relaxations; hierarchical models; Rao-Blackwellization; iterative proportional fitting (IPF); maximum-entropy (MAXENT) models and iterative scaling; L1-regularization and logistic regression

- **Decision-theoretic Graphical Models**

- *Framework*: utility theory; decision theory

- *Models*: influence diagrams (aka decision networks)
- *Concepts*: decision-making under uncertainty; preferences; maximum-expected-utility principle
- *Problems*: optimal action selection
- *Advanced topics*: factored (partially observable) Markov decision processes ((PO)MDPs); policy selection; decentralized (PO)MPDs; multi-agent coordination; (factored) value and policy iteration; basis-function approximations; sparse sampling

Advanced Topics

• Stochastic Graphical Models

- *Framework*: Stochastic processes
- *Models*: factored hidden Markov models (fHMMs); temporal/dynamic Bayesian networks (DBNs)
- *Problems*: Viterbi/most-likely sequence; belief updating; belief-state tracking/estimation; prediction/forecasting
- *Methods*: Viterbi and forward-backward algorithms; factored belief-state projections; particle filter (PF), aka survival-of-the-fittest
- *Learning from data*: Baum-Welch algorithm

• Game-theoretic Graphical Models

- *Framework*: Noncooperative game theory
- *Model*: Graphical games
- *Problems*: computing Nash and correlated equilibria
- *Concepts*: CSP formulation; graphical potential games and connections to probabilistic inference; decomposable graphical models
- *Methods*: NashProp; heuristic local search (e.g., hill climbing); linear feasibility systems and LPs
- *Advanced topics*: mathematical (competitive) economics; stochastic games; multi-agent influence diagrams (MAIDs); graphical economies; market equilibrium; continuation methods; ADProp; Papadimitriou's ellipsoid-based algorithm; no-regret learning; regret matching

• Causal, Econometric and Behavioral Graphical Models

- *Framework*: Cognitive Science and Psychology; Econometrics; Social Network Theory
- *Models*: Causal networks; simultaneous equation models (SEMs); linear threshold models (LTMs); generalized threshold models (GTMs);
- *Problems*: causal reasoning and prediction; identifying the most “influential” nodes; learning from behavioral data

- **Graphical Models for First-order Logic**

- *Models*: probabilistic relational models (PRMs); relational Markov networks; logical HMMs; dynamic probabilistic relational models (DPRMs); logical Markov decision processes

Motivating applications in computer vision, speech recognition, artificial intelligence (e.g., reasoning and decision-making under uncertainty) and machine learning will be embedded within the presentation of the different topics

NOTE: *The list of topics, as well as the emphasis on each topic, will likely vary depending on the background and interests of the course participants.*

Assessment

Course Project: Students complete a research project on a relevant topic. The specific problem requires the instructor's approval. Possible projects include, but are not limited to, an application of a particular model and technique to a specific problem, the development of a system based on graphical model ideas to solve a particular problem for a particular application domain, or novel experimental evaluations and comparisons of one or various computational techniques. Ideally, the project would incorporate a combination of theoretical and experimental work. Project that involve exclusively theoretical work may be permitted only under very close consultation with the instructor and should involve a narrowly and clearly defined problem description and line-of-attack. Students periodically submit progress reports to monitor the project's development. Students give a presentation and submit a written report on the results of the project by the end of the course.

Student Evaluations: Students performance will be evaluated based on their level of participation during the discussions, and the quality of their topic oral presentation and written report, as well as their project's proposal, progress reports, oral presentation and final written report. There may also be sporadic written and/or code/implementation homework assignments (at the discretion of the instructor).

Grades: The following table shows the amount and weighting of each evaluation component in the course.

Criteria	Quantity	Percent
Class participation		20%
Homework		10%
Topic report	1	10%
Topic oral presentation	1	10%
Project proposal	1	5%
Project progress reports	2	5%
Project oral presentation	1	15%
Project final report	1	25%
Total		100%

Final grades will be assigned using the traditional grading scale (90-100=A, 80-89=B, 70-79=C, 60-69=D, 0-59=F), with deviations at the instructor's discretion.

General University Statements

Americans with Disabilities Act: If you have a physical, psychological, medical or learning disability that may impact your course work, please contact Disability Support Services, ECC (Educational Communications Center) Building, room128, (631) 632-6748. They will determine with you what accommodations, if any, are necessary and appropriate. All information and documentation is confidential.

Academic Integrity: Each student must pursue his or her academic goals honestly and be personally accountable for all submitted work. Representing another person's work as your own is always wrong. Faculty are required to report any suspected instances of academic dishonesty to the Academic Judiciary. Faculty in the Health Sciences Center (School of Health Technology & Management, Nursing, Social Welfare, Dental Medicine) and School of Medicine are required to follow their school-specific procedures. For more comprehensive information on academic integrity, including categories of academic dishonesty, please refer to the academic judiciary website at <http://www.stonybrook.edu/uaa/academicjudiciary/>

Critical Incident Management: Stony Brook University expects students to respect the rights, privileges, and property of other people. Faculty are required to report to the Office of Judicial Affairs any disruptive behavior that interrupts their ability to teach, compromises the safety of the learning environment, or inhibits students' ability to learn. Faculty in the HSC Schools and the School of Medicine are required to follow their school-specific procedures.