Rutgers University
Department of Electrical and Computer Engineering, Spring 2024

Probabilistic Graphical Models (PGM) & Inference Algorithms (Graduate Course)

Probabilistic Graphical Models (PGM) are based on graph, probability, estimation, and information theory, as well as elements of machine learning and offer a fascinating unifying theoretical framework exploited in a rich variety of (challenging) engineering applications, including **communications, computer vision, natural language processing, bioinformatics, social networks, and big data analysis**. They are particularly useful in problems that can be described as graphs of random variables, and their theory is currently an active topic of research. More specifically, PGMs encode (conditional) dependencies among random variables on carefully crafted graphs. Such description is powerful enough to describe a variety of many famous algorithms, such as (Gaussian) Belief Propagation, Kalman Filtering, Viterbi, Expectation-Maximization (EM). This class will offer an introduction in representation with PGMs, algorithms for exact inference, approximate inference, and learning/estimation.
More specifically: Directed acyclic graphs (Bayesian Nets) factorization theorem and semantics (I-map, d-separation, P-map). Undirected graphs (Markov Blanket, Hammersley-Clifford theorem), factor graphs (and techniques to convert), Gaussian Graphical Models. Exact Inference (elimination algorithm, sum-product/belief propagation, max-product on Trees, HMMs and Kalman Filtering, Junction Tree algorithm). Approximate Inference: Loopy Belief Propagation, Sampling Methods (Particle Filtering, Metropolis-Hastings). Intro to learning graphs: ML Techniques, Chow-Liu, BIC-based Techniques, EM.
Term projects will thoroughly study application examples in diverse domains.

Graduate class 16:332:539:03/14:332:436:08, credits: 3, also available to senior undergraduates under Permission of Instructor.

Instructor: Prof. Aggelos Bletsas

Lectures: Monday, Thursday, 12:10 - 1:30 @ CoRE 538 (Busch campus).
Office hours: CoRE 530, Monday 13.30-13.30.

Grading: 30% Midterm and Psets, 30% Term Project, 40% Final.

Books and Course Notes: The syllabus will be based on MIT 6.438 notes, Koller and Friedman textbook, papers, as well as selected material:

D. Koller and N. Friedman, Probabilistic Graphical Models: Principles and Techniques, MIT Press, 2009.

C. M. Bishop, Pattern Recognition and Machine Learning, Springer Verlag, 2006.

D. Barber, Bayesian Reasoning and Machine Learning, Cambridge University Press, 2012.

MIT “Algorithms for Inference” 2014 graduate course notes (also given as 6.438), MITOPENCOURSEWARE (OCW), available at
https://ocw.mit.edu/courses/6-438-algorithms-for-inference-fall-2014/pages/lecture-notes/

Tentative schedule (similar to MIT 6.438 Inference class)

Lecture #	Topic
1	Introduction, overview, logistics, probability theory review
2	Directed Acyclic Graphs (DAGs) (Bayesian Nets) & Factorization Theorem
3	Undirected Graphical Models Relation to DAGs, MRFs & Markov Blanket, Hammersley-Clifford Theorem
4	Factor Graphs Techniques for converting one PGM representation to the other
5-6	Minimal I-Maps, Chordal Graphs, Trees, and Markov Chains (Notions of I-map and P-map)
7	Gaussian Graphical Models Jointly Gaussian Random Variables, Operation on Gaussian Vectors, Gaussian PGMs and Matrix Inversion Lemma
8	Exact Inference: Elimination Algorithm
9	Treewidth & Elements of Graph Theory
10	Exact Inference: Sum-Product on Trees, Parallel Sum-Product
11	Exact Inference: Forward-Backward Algorithm (Hidden Markov Models (HMMs)), Example on Convolutional Decoding
12	Exact Inference: Sum-product on factor tree graphs, MAP elimination algorithm
13-14	Exact Inference: Max-Product on Trees Max-sum, Min-sum variations, Example on Convolutional Decoding
	Midterm
15	Gaussian Belief Propagation (BP)
16	BP on Gaussian HMMs: Kalman Filtering
17	Junction Tree Algorithm
18-19	Approximate Inference: Loopy Belief Propagation
20-21	Intro to Learning Graphical Models: ML method, Bayesian Parameter estimation, hyperparameters
22	Intro to Learning Graphical Models: Learning parameters of an undirected graphical model
23-24	Intro to Learning Graphical Models: Learning Structure in Directed Graphs Chow-Liu algorithm, Bayesian Score
25	Markov Chain Monte Carlo (Sampling) Methods and Approximate MAP Metropolis-Hastings and mixing time
26	Approximate Inference: Importance Sampling, Particle Filters
27	Parameter Estimation from Partial Observations: Expectation-Maximization Algorithm Term Project Presentations
	Final (Thursday May 2)