EE512 Graphical Models
Spring 2006

University of Washington
Dept. of Electrical Engineering

Handout 1: Course Overview

Prof: J. Bilmes <bilmes@ee.washington.edu>

March 28^{th}, 2006

Description: This course will serve as a thorough overview of graphical models including Bayesian networks, undirected models, and factor graphs. Topics will include the following, which also serves as a rough weekly outline (i.e., if we stay on schedule, the number is also the week number in the quarter).

1) Intro. Graph types: conditional independence; directed, undirected, and factor models; algorithms for conditional independence (e.g., Bayes-ball, d-separation, Markov properties on graphs, factorization, Hammersley-Clifford theorems).

2) Models: linear Gaussian models, mixture models, factor analysis, probabilistic decision trees, Markov Random Fields, Gibbs distributions, static conditional random fields (CRFs), multivariate Gaussians as graphical models.

3) Dynamic (temporal) models: Hidden Markov Models, Kalman filtering and linear-Gaussian HMMs, linear dynamic models, dynamic Bayesian networks (DBNs), label and observation bias in natural language processing, dynamic conditional random fields (CRFs), and general dynamic graphical models.

4) Chordal Graph Theory: moralization; triangulated, decomposable, and intersection graphs, k-trees, hyper-graphs. Tree-width and path-width parameters of a graph. Graph separators and their crossings. Relational data-base systems and joins, and phylogenetic tree construction.

5) Exact Probabilistic Inference: junction trees, belief propagation in its various forms (including Pearl's formulation, Hugin, Shafer-Shenoy, Bucket-elimination, etc.); join-trees and data-structures; optimal triangulations. The elimination family of algorithms. Generalized distributed law (GDL). Relation to dynamic programming. Generality (such as Viterbi, MPE, the fast Fourier transform). NP hardness results, and NP hardness results for finding the best form.

6) Inference as search: Search-based approaches to inference and how it relates to junction trees, message passing, and BP. Results from the constraint satisfaction (CSP) and SAT literatures, constraint solvers, and heuristics used in these domains. Cutset and recursive conditioning, comparison with elimination. Branch-and-bound.

7) Approximate Probabilistic Inference: loopy belief propagation (BP), expectation propagation (EP), generalized belief propagation (GBP), cluster GBP, comparisons with mean-field and variational approaches, statistical physics, Bethe and Kikuchi free-energies, fixed points, recent theoretical results and open questions. Also, sampling approaches such as MCMC and particle algorithms (such as Rao-Blackwellized particle filtering). Pruning based approaches.

8) Learning: learning Bayesian networks, hardness results for learning networks in the maximum likelihood sense, learning in the PAC (probably approximately correct) framework, EM for parameter and structure learning, alternating minimization, discriminative parameter and structure learning, CRF training vs. discriminatively trained HMM, other learning methods. Information theory and graphical models.

9) Models in practice: Real-world static graphs. HMMs and DBNs in speech, language, and bioinformatics, QMR and Factorial HMMs, turbo-coding, low-density parity check codes, other codes on graphs, belief propagation algorithms on these graphs. Practical issues such as data structures and algorithms useful for performing inference.

The course will have homework assignments (about one per week or so), a midterm exam (which will be take home), and a final project (which will require project status reports, a final research paper, and a final presentation). There will be no final exam.

Homeworks will be due on Fridays, at 5:00pm — either print out your solutions and place them in the inbox outside my door (EE1-418), or email them in Acrobat .pdf format to **both** the TA and myself.

We have 20 lectures in total. Since the final project will be the last day of lecture, there are 19 lecture days total. There will roughly be 2 lectures per topic above, which leaves us one "buffer" day.

Lecturer: Prof. Jeff Bilmes 418 EE1 Bldg. `bilmes@ee.washington.edu` Lecture sessions: Tu/Thu 1:30-3:30 EE1-054. Office Hours: TBA

Ta: Chris Bartels, 203 EE1 Bldg. `bartels@ee.washington.edu` Discussion sessions: TBA. Office Hours: Wednesdays, 3:30-4:30, EE Student Center, Sieg 127.

Course web page: `http://ssli.ee.washington.edu/courses/ee512`

Text: The main text will be "An Introduction to Graphical Models", by Mike Jordan. This text has not yet been published, but we will use individual chapters from the text to be available in the copy center `b042cmu@u.washington.edu`. Please email me any suggestions, typos, etc. that you might find in the text. We will also read chapters from "Graphical Models" by S. Lauritzen.

Class Reader: We will also use an online course "reader" consisting of slides, possible class handouts and also papers generally available on the web.

Prerequisites: Statistical Pattern Recognition I, prior exposure to pattern recognition concepts, or consent of instructor. Also, you will need to know some programming language (such as C, C++, Java, or matlab, any language is fine). The course is open to students in all UW departments.

Grading: Grades will be based on a combination of the final project (35%), homeworks (35%), and the midterm exam (30%). There will be a new home due everyone 1-2 weeks.

A P under P/NP grading (or S/NS) requires doing the homeworks and the final project (but not the midterm). We have existing scribes from the previous time this course was offered that you might find useful to look at.

Project: A significant portion of your grade will be based on the class project that involves using graphical models in some way. There are a number of tools that are available for doing this, such as BNT or GMTK. Projects can be done in groups of no more than two people.

Other Books you might find useful.

- "An Introduction to Bayesian Networks", F.V. Jensen, 1996. A good general introduction to Bayesian networks (out of print, but available in the library).
- "Bayesian Networks and Decision Graphs", F.V. Jensen, 2001. Another good general introduction to Bayesian networks (not out of print).
- "Graphical Models", S.L. Lauritzen. Oxford, 1996. A very complete, theoretically precise, but dense text, authored by one of the field's leading authorities.
- "Probabilistic Networks and Expert Systems", R.G. Cowell, A.P. Dawid, S.L. Lauritzen, and D.J. Spiegelhalter. 1999. Similar to the previous text, but includes more material on inference, applications, and other general problems.
- "Artificial Intelligence: A Modern Approach: 2nd Edition", S. Russel and P. Norvig, 2003. Has a nice introductory chapter on Bayesian networks.
- "Learning in Graphical Models", Ed. by M.I. Jordan. An excellent collection of recent research papers compiled by Mike Jordan, one of the leading experts in this field.
- "Probabilistic Reasoning in Intelligent Systems", J. Pearl. 1998. A classic early text by one of the founders of the field. Pearl is credited with inventing the term "Bayesian networks".
- "Causality", J. Pearl. 2000. A relatively newer text by Pearl specifically on causality, and causal modeling.
- "Pattern Classification", R. Duda, P. Hart and D. Stork (the text used for 596I). The original text (published in 1973) is still widely read.

- "The Elements of Statistical Learning: Data Mining, Inference, and Prediction", Hastie, Tibshirani, and Friedman.
- "Neural Networks for Pattern Recognition", by C. Bishop, 1996. (available now in the UW bookstore). This book mainly contains background material, but has become a classic text in pattern recognition even though "neural networks" is in the title, and is worth reading if you plan to do any work at all in pattern recognition.

Important Dates:

Holiday: Monday, May 29th, Memorial day (but doesn't influence us).

Take-Home Midterm: Out, May 5th. Due Friday, May 12th.

Final Project Presentations: In class final presentations will during the date of last lecture: Thursday, June 1st, EE1-054, 1:30-4:30 (we will use an extra hour).

Final Project Report: The 4-page final project report will be due **electronically**, Tuesday June 6th, 2006.