Part A

Advanced Topics in Bayesian Networks

Team members:
  Tom Temple
  Ethan Howe
  James Lenfestey

Topics:
  Brief Intro to Bayes Networks
  Dynamic Bayes Networks
    Exact inference (is intractable)
    Approximate Inference (PF)
    Learning (Parameter/Structure estimation)

  Probabilistic Relational Models
    Background: statistical models of real-world relational data
    Can be flattened to BNs (richer but no more general)
    Parameter/Structure estimation

Part B

Learning Dynamic Bayesian Networks. Ghahramani.
ftp://ftp.cs.toronto.edu/pub/zoubin/vietri.ps.gz

Learning Probabilistic Relational Models. Friedman, Getoor, Koller, Pfeffer
http://ai.stanford.edu/~koller/papers/ijcai99lprm.ps

Part C

In our lecture we will discuss extensions of Bayesian models that are designed
to deal with temporal and relational structures. These models address Bayesian

models that otherwise have a high degree of regularity. Temporal networks have fixed models of transition dynamics and flattened relational data has template structures that would otherwise appear as isomorphic subgraphs. As a way to deal with temporal regularity, we introduce Dynamic Bayesian Networks, which are a concise alternative to Hidden Markov Models. Probabilistic relational models summarize relationships between groups of variables. Each method is important to cognitive robots because robots must both deal with an evolving uncertain state and classes of similar objects and situations. Each has myriad applications, e.g. DBNs have been used for speech recognition, and PRMs for data mining.

Part D

Artifical Intelligence, A Modern Approach. Russell and Norvig.
Chapter 14: Intro to Bayesian networks. Probabilistic inference. PRM primer.
Chapter 15: Temporal Bayesian models. HMMs and DBNs.

Learning Bayesian Networks from Incomplete Data with Stochastic Search Algorithms. Myers, Laskey, Levitt
ite.gmu.edu/~klaskey/papers/myersUAI99.pdf
This paper discusses learning in general Bayesian networks. It describes evolutionary and Markov Chain Monte Carlo algorithms for doing approximate learning, and then combines the two approaches into an improved algorithm.

Dynamic Probabilistic Relational Models. Sanghai, Domingos, Weld.
http://www.cs.washington.edu/homes/weld/papers/sanghai-ijcai03.pdf
DPRMs are the dynamic version of PRMs where each time-slice and its dependence on the previous slice is modelled with a PRM. Although we won't cover DPRMs they are interesting as the natural junction of the two concepts that we will present.

Speech Recognition with Dynamic Bayesian Networks
www.cs.berkeley.edu/~russell/papers/aaai98-speech.ps
Speech recognition is a successful application of inference in DBNs. This paper presents a DBN model that is more successful than comparable hidden Markov models.

Part E

3 part lecture:

I. Intro to BNs. Tom Temple.

II. DBNs.
  Definition. Tom Temple
  Inference/Learning. Ethan Howe.
  Demo. Ethan Howe.

III. PRMs. James Lenfestey

Part F

We plan to construct a simple DBN and perform inference subject to observed data. To do this we will use the Bayes Net Toolbox for Matlab written by Kevin Murphy (http://www.cs.ubc.ca/~murphyk/Software/BNT/bnt.html). The package provides a graphical interface, so we can construct the DBN with the help of class participation.