Bayesian AI Tutorial

(A 3 HOUR TUTORIAL WITH OPTIONAL LAB)

http://www.csse.monash.edu.au/bai/

Kevin B. Korb and Ann E. Nicholson
School of Computer Science and Software Engineering
Monash University, Victoria 3800

{korb,annn}@csse.monash.edu.au
Phone: (3) 9905-5198/5211
Fax: (3) 9905-5146

Bayesian networks have rapidly become one of the leading technologies for reasoning under uncertainty, by explicitly modeling causal relationships as well as supporting decision analysis. This follows the work of Pearl, Lauritzen, and others in the late 1980s using graphical models to make Bayesian reasoning feasible.

In this tutorial, we begin with a brief examination of the philosophy of Bayesianism, motivating the use of probabilities in decision making, agent modeling and data analysis. We introduce Bayesian networks for modeling and reasoning under uncertainty and provide an overview of the inference techniques, including causal inference. We then describe two extensions of Bayesian networks: (1) decision networks, which explicitly support decision making under uncertainty, and (2) dynamic Bayesian networks, which allow explicit reasoning about changes over time. We will illustrate the presentation throughout with examples using the Netica Bayesian network software.

The biggest obstacle to Bayesian AI having a broad and deep impact outside of the research community are the difficulties in developing applications, difficulties with eliciting knowledge from experts, and integrating and validating the results. This has led to the rapid growth in applying machine learning methods to the automated building of Bayesian networks. In this tutorial, we review the main data mining techniques available for learning Bayesian network structures and for parameterizing them, including dealing with incomplete data.

Another issue is that there is as yet no clear methodology for developing, testing and deploying Bayesian network technology in industry and government — there is no recognized discipline of "software engineering" for Bayesian networks. In this tutorial we introduce the process of Knowledge Engineering with Bayesian Networks (KEBN), together with some tools we have developed in support.

We conclude the presentation part of the tutorial with two case studies of Bayesian network development: the first an ecological risk assessment model developed in the Monash Water Studies centre, and the second SARBayes, a collaborative project with Victorian Search and Rescue. Following the formal tutorial presentation in the morning, there will be an afternoon lab session where participants will undertake some modeling exercises using Netica.

Outline

1. Bayesian AI
   (a) Probabilities
   (b) Bayesian philosophy
   (c) Probabilistic Causality
   (d) Decision making and AI
2. Introduction to Bayesian networks
   (a) Nodes, structure and probabilities
   (b) Reasoning with Bayesian networks
   (c) Understanding Bayesian networks
3. Inference
   (a) Exact algorithms
   (b) Approximate algorithms
   (c) Causal modeling
4. Extensions
   (a) Decision networks
   (b) Dynamic Bayesian networks

5.. Causal Discovery
   (a) Learning Bayesian network structure
   (b) Parameterizing Bayesian networks
   (c) Incomplete data
   (d) MCMC learning methods

6. Knowledge Engineering of Bayesian Networks (KEBN)
   (a) The KEBN process
   (b) Tools for supporting KEBN
7. Example applications
   (a) Ecological risk assessment (Carmel Pollino, Water Studies)
   (b) SARBayes: Search and Rescue
8. Laboratory Session
   (a) Introduction to Netica
   (b) Exercise 1: Building a Bayesian network
   (c) Exercise 2: Parameterizing Bayesian networks
   (d) Exercise 3: Causal discovery

Presenter Biographies

Ann Nicholson completed her doctorate in the robotics research group at Oxford University (1992) working on dynamic belief networks for discrete monitoring. She then spent two years at Brown University as a post-doctoral research fellow before taking up a lecturing position at Monash University in Computer Science. Her general research focus is AI methods for reasoning under uncertainty, while her current research areas include approximate methods for Bayesian network evaluation, applications of dynamic belief networks and approximate planning using Markov Decision Processes. She has taught numerous subjects in Computer Science, including Artificial Intelligence, and developed and taught a web-based introduction to Lisp programming.

Kevin Korb did his PhD in the philosophy of science at Indiana University (1992) working on the philosophical foundations for the automation of Bayesian reasoning. Since then he has lectured at Monash University in Computer Science, combining his interests in philosophy of science and artificial intelligence in work on understanding and automating inductive inference, the use of MML in learning causal theories, artificial evolution of cognitive and social behavior, and modeling Bayesian and human reasoning in the automation of argumentation. He has independently developed and taught the following subjects: Machine Learning (3rd year and honours); Bayesian Reasoning (honours); Causal Reasoning (honours); The Computer Industry: Historical, Social and Professional Issues (3rd year).