Probabilistic Reasoning in AI

308-526

Instructor: Doina Precup

Email: dprecup@cs.mcgill.ca

TA: Will Renner

Email: wrenner@acm.org

Class web page

http://www.cs.mcgill.ca/∼dprecup/courses/Winter2002/prob.html

Outline

- Administrative details
- Dealing with uncertainty
- Probability
- Probabilistic reasoning
- Decision making under uncertainty

Administrative issues

- Class material:
  - J. Pearl, Probabilistic Reasoning in Intelligent Systems
  - R.S.Sutton and A.G.Barto,
    Reinforcement Learning: An Introduction
  - Class notes: posed on the web page
  - Additional readings: TBA
- Evaluation mechanisms:
  - Six problem sets (40%)
  - Four programming assignments (20%)
  - Two reading assignments (10%)
  - Class participation and discussions (up to 5% extra credit)
- Programming assignments must **function** to get credit

Uncertainty

Uncertainty is inherent in many tasks

E.g. Will leaving home t minutes before the flight get me to the airport on time?

- Partial knowledge of the state of the world

  E.g. We do not know the road state, other drivers’ plans etc.

- Noisy observations

  E.g. Traffic reports

- Inherent stochasticity

  E.g. Flat tires, accidents etc.

- Phenomena that are not covered by our models

  E.g. the complexity of predicting traffic

How do we deal with uncertainty?

- Implicit methods
  - Ignore uncertainty as much as possible
  - Build procedures that are robust to uncertainty
  - E.g. AI planning methods
- Explicit (model-based) methods
  - Build a *model of the world* that describes the uncertainty about the system state, dynamics and about our observations
  - Reason about the effect of actions given the model

We will focus mainly on explicit model-based methods

How do we represent uncertainty?

- What language should we use? What are the semantics of our representations?
- What queries can we answer with our representations? How do we answer them?
- How do we construct a representation? Do we need to ask an expert, or can we learn from data?

Why logic breaks

A purely logical approach either:

1. risks falsehood

   E.g. leaving 25 minutes early will get me to the airport on time

2. leads to conclusions that are too weak for decision making:

   E.g. Leaving 25 minutes early will get me to the airport on time if there is no accident on the bridge and it does not rain and my tires remain intact etc. etc.”

Probability

- A well-known and well-understood framework for dealing with uncertainty
- Has a clear semantics
- Provides principled answers for:
  - Combining evidence
  - Predictive and diagnostic reasoning
  - Incorporation of new evidence
- Can be learned from data
- Arguably intuitive to human experts

Representing probabilities efficiently

- Naive representations of probability are hopelessly inefficient

  E.g. consider patients described by several attributes:
  - background: age, gender, medical history,...
  - Symptoms: fever, blood pressure, headache,...
  - Diseases: pneumonia, hepatitis,...
- A probability distribution needs to assign a number to each combination of values of these attributes!
- Real examples involve hundreds of attributes
- **Key idea:** exploit regularities and structure of the domain
- We will focus mainly on exploiting **conditional independence** properties

A Bayesian network

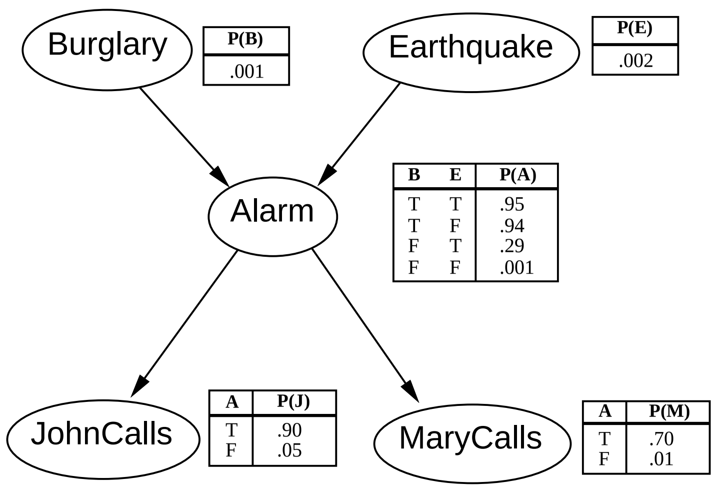

P(B)
.001

P(E)
.002

B	E	P(A)
T	T	.95
T	F	.94
F	T	.29
F	F	.001

A	P(J)
T	.90
F	.05

A	P(M)
T	.70
F	.01

Probabilistic Reasoning

During the first half of the course we will study:

- Syntax and semantics of Bayesian networks
- How to efficiently answer queries in a Bayesian network
- How to learn Bayesian networks from data
- How to extend Bayesian networks in order to represent properties of sequences and temporal processes

Decision making

- Probability is not enough for choosing actions
- We also need to consider *risks and payoffs*

E.g. Suppose I believe the following:

$$\begin{aligned}
P(A_{25} \text{ gets me there on time} | \ldots) &= 0.04 \\
P(A_{90} \text{ gets me there on time} | \ldots) &= 0.70 \\
P(A_{120} \text{ gets me there on time} | \ldots) &= 0.95 \\
P(A_{1440} \text{ gets me there on time} | \ldots) &= 0.9999
\end{aligned}$$

Which action should I choose?

Depends on my *preferences* for missing flight vs. airport cuisine, etc.

Utility theory is used to represent and infer preferences

Decision theory = utility theory + probability theory

Practical decision making

- We need to represent both probabilities and utilities
- The **expected utility** of actions is computed given evidence and past actions
- We choose the action that maximizes expected utility
- **Value of information:** is ti worth acquiring more information in order to choose better actions?

Decision making

In the second half of the course we will study:

- Utility theory
- Models of repeated decision: Markov Decision Processes
- Partially Observable Markov Decision Processes
- Learning to act optimally

Related fields

- Artificial Intelligence
- Machine learning
- Operations research
- Decision theory
- Statistics
- Information theory
- ...

Fielded applications

- Expert systems
  - Medical diagnosis (e.g. Pathfinder)
  - Fault diagnosis (e.g. jet-engines)
- Monitoring
  - Space shuttle engines (Vista project)
  - Freeway traffic
- Sequence analysis and classification
  - Speech recognition
  - Biological sequences
- Information access
  - Collaborative filtering
  - Information retrieval

Example: Pathfinder (Heckerman, 1991)

- Medical diagnostic system for lymph node diseases
- Large net! 60 diseases, 100 symptoms and test results, 14000 probabilities
- Network built by medical experts
  - 8 hours to determine the variables
  - 35 hours for network topology
  - 40 hours for probability table values
- Experts found it easy to invent causal links and probabilities
- Pathfinder is now **outperforming world experts** in diagnosis
- Commercialized by Intellipath and Chapman Hall Publishing; being extended now to other medical domains