Decision Networks

Yuqing Tang

Doctoral Program in Computer Science
The Graduate Center
City University of New York
ytang@cs.gc.cuny.edu

Outline

Introduction

- Bayesian networks can be extended to support decision making.
- Preferences between different outcomes of various plans.
  - Utility theory
- Decision theory = Utility theory + Probability theory.

Bayesian Decision Theory

- Frank Ramsey (1926)
  Decision making under uncertainty: what action to take (plan to adopt) when future state of the world is not known.

Bayesian answer: Find utility of each possible outcome (action-state pair) and take the action that maximizes expected utility.

Example

action	Rain ($p = 0.4$)	Shine ($1 - p = 0.6$)
Take umbrella	30	10
Leave umbrella	-100	50

Expected utilities:

$$\begin{aligned} E(\textit{Take umbrella}) &= (30)(0.4) + (10)(0.6) = 18 \\ E(\textit{Leave umbrella}) &= (-100)(0.4) + (50)(0.6) = -10 \end{aligned}$$

Outline

Decision Networks

A Decision network represents information about

- the agent's current state
- its possible actions
- the state that will result from the agent's action
- the utility of that state

Also called, Influence Diagrams (Howard&Matheson,1981).

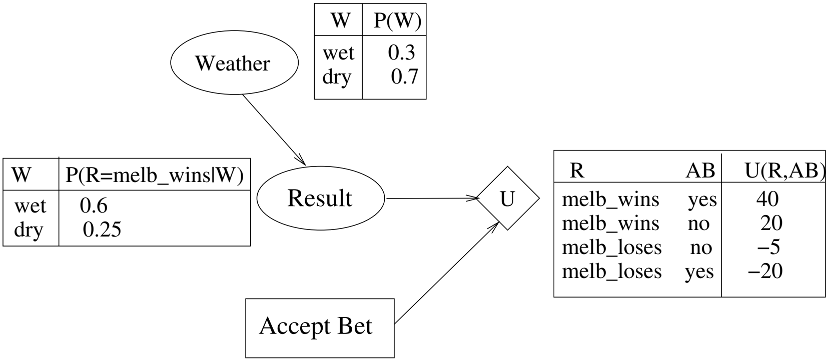

W	P(W)
wet	0.3
dry	0.7

W	P(R=melb_wins|W)
wet	0.6
dry	0.25

R	AB	U(R,AB)
melb_wins	yes	40
melb_wins	no	20
melb_loses	no	–5
melb_loses	yes	–20

Type of nodes

- Chance nodes: (ovals) represent random variables (same as Bayesian networks). Has an associated CPT. Parents can be decision nodes and other chance nodes.
- Decision nodes: (rectangles) represent points where the decision maker has a choice of actions.
- Utility nodes: (diamonds) represent the agent's utility function (also called value nodes in the literature). Parents are variables describing the outcome state that directly affect utility. Has an associated table representing multi-attribute utility function.

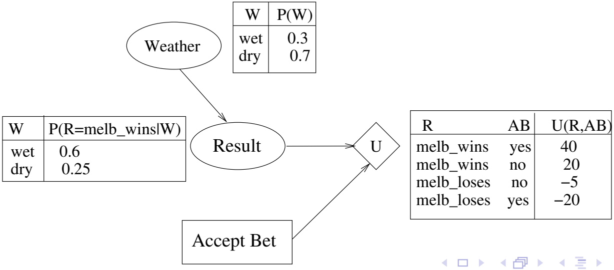

W	P(W)
wet	0.3
dry	0.7

W	P(R=melb_wins|W)
wet	0.6
dry	0.25

R	AB	U(R,AB)
melb_wins	yes	40
melb_wins	no	20
melb_loses	no	−5
melb_loses	yes	−20

Decision networks (example)

Example

Clare's football team, Melbourne, is going to play her friend John's team, Carlton. John offers Clare a friendly bet: whoever's team loses will buy the wine next time they go out for dinner. They never spend more than $15 on wine when they eat out. When deciding whether to accept this bet, Clare will have to assess her team's chances of winning (which will vary according to the weather on the day). She also knows that she will be happy if her team wins and miserable if her team loses, regardless of the bet.

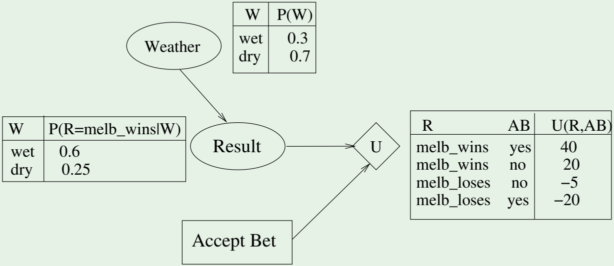

W	P(W)
wet	0.3
dry	0.7

W	P(R=melb_wins|W)
wet	0.6
dry	0.25

R	AB	U(R,AB)
melb_wins	yes	40
melb_wins	no	20
melb_loses	no	–5
melb_loses	yes	–20

Expectation and expected utilities

Expectation

$$E(X) = \Sigma_{v \in Domain(X)} v \cdot P(X = v)$$

Expected utility of an action given evidence

$$EU(A|\mathbf{E}) = \Sigma_i P(O_i|\mathbf{E}, A) U(O_i|A)$$

- **E** is the available evidence
- A is an action taken
- O_i is one of the possible outcome state
- U is the utility function which measures the utility of the outcome O_i given the action A

Evaluating Decision Networks

1. Add any available evidence.
2. For each action value in the decision node:
    1. Set the decision node to that value;
    2. Calculate the posterior probabilities for the parent nodes of the utility node, as for Bayesian networks, using a standard inference algorithm;
    3. Calculate the resulting expected utility for the action.
    4. Return the action with the highest expected utility.

Evaluating Decision Networks: Example

$$
\begin{aligned}
P(R = melb\_wins) &= P(W = w)P(R = melb\_wins|W = w) \\
&\quad + P(W = d)P(R = melb\_wins|W = d) \\
EU(AB = yes) &= P(R = wins)U(R = wins|AB = yes) \\
&\quad + P(W = loses)P(R = loses|AB = yes) \\
&= (0.3 \times 0.6 + 0.7 \times 0.25) \times 40 \\
&\quad + (0.3 \times 0.4 + 0.7 \times 0.75) \times (-20) \\
&= 0.355 \times 40 + 0.645 \times (-20) = 14.2 - 12.9 \\
&= 1.3 \\
EU(AB = no) &= P(R = wins)U(R = wins|AB = no) \\
&\quad + P(W = loses)P(R = loses|AB = no) \\
&= (0.3 \times 0.6 + 0.7 \times 0.25) \times 20 \\
&\quad + (0.3 \times 0.4 + 0.7 \times 0.75) \times (-5) \\
&= 0.355 \times 20 + 0.645 \times (-5) = 7.1 - 3.225 \\
&= 3.875
\end{aligned}
$$

Information Links

- Indicate when a chance node needs to be observed before a decision is made.

W	F	P(F|W)
wet	rainy	0.60
	cloudy	0.25
	sunny	0.15
dry	rainy	0.10
	cloudy	0.40
	sunny	0.50

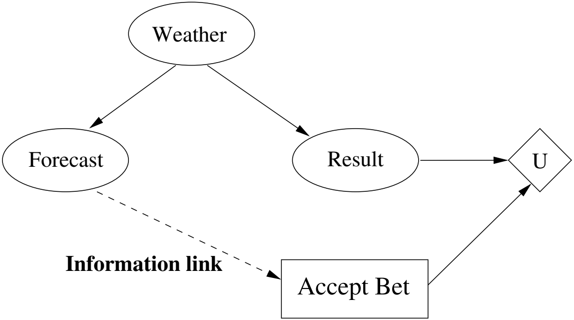

Decision Table

F	Accept Bet
rainy	yes
cloudy	no
sunny	no

Decision Table Algorithm

1. Add any available evidence.
2. For each combination of values of the parents of the decision node:
    1. For each action value in the decision node:
        1. Set the decision node to that value;
        2. Calculate the posterior probabilities for the parent nodes of the utility node, as for Bayesian networks, using a standard inference algorithm;
        3. Calculate the resulting expected utility for the action.
    2. Record the action with the highest expected utility in the decision table.
3. Return the decision table.

Outline

1. Introduction
2. Decision networks
   - Decision networks with intervening actions
   - Dynamic Belief Networks
3. Summary

Fever problem description

Example

Suppose that you know that a fever can be caused by the flu. You can use a thermometer, which is fairly reliable, to test whether or not you have a fever. Suppose you also know that if you take aspirin it will almost certainly lower a fever to normal. Some people (about 5% of the population) have a negative reaction to aspirin. You'll be happy to get rid of your fever, as long as you don't suffer an adverse reaction if you take aspirin.

Fever decision network

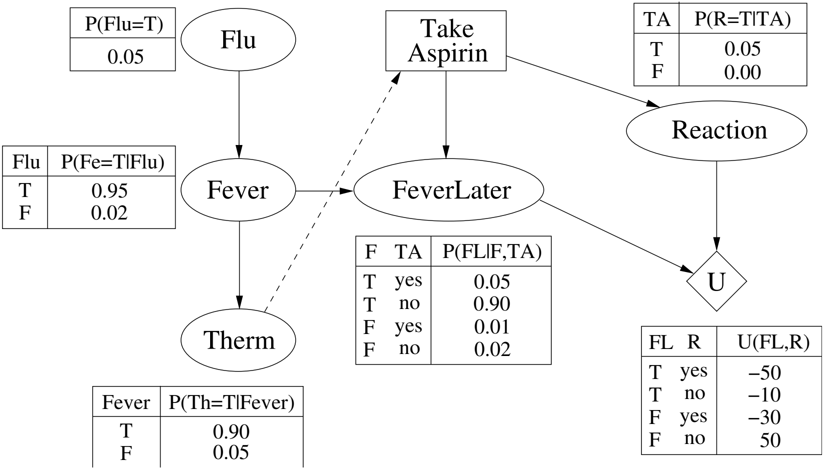

P(Flu=T)
0.05

Flu	P(Fe=T|Flu)
T	0.95
F	0.02

Fever	P(Th=T|Fever)
T	0.90
F	0.05

TA	P(R=T|TA)
T	0.05
F	0.00

F	TA	P(FL|F,TA)
T	yes	0.05
T	no	0.90
F	yes	0.01
F	no	0.02

FL	R	U(FL,R)
T	yes	–50
T	no	–10
F	yes	–30
F	no	50

Fever decision table

Evidence	$Bel(FLater = T)$	$EU(TA = yes)$	$EU(TA = no)$	Decision
None	0.046	45.27	45.29	no
Th=F	0.525	45.41	48.41	no
Th=T	0.273	44.1	19.13	yes
Th=T & Reation=T	0.273	-30.32	0	no

Types of actions

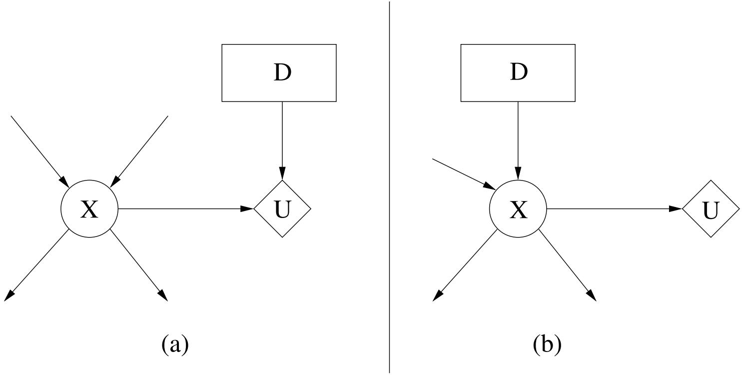

(a) Non-intervening and (b) Intervening

Outline

Sequential decision making

- Precedence links used to show temporal ordering.
- Network for a test-action decision sequence

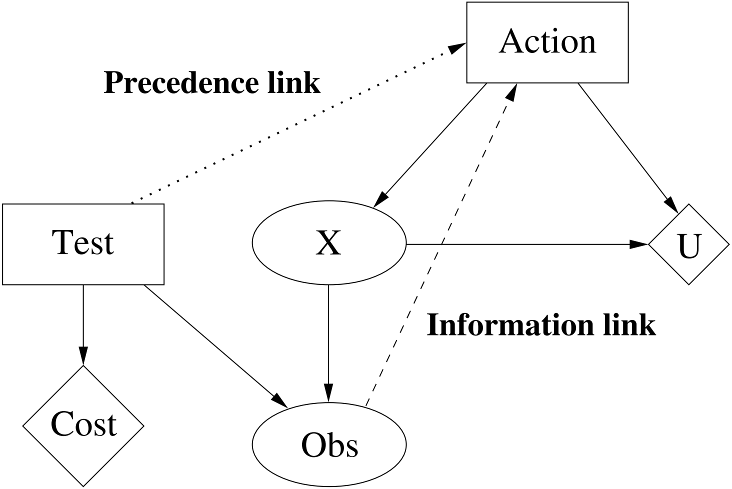

Dynamic Bayesian networks

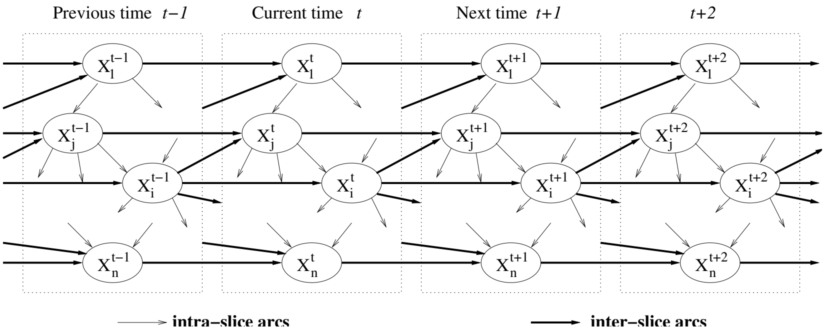

- One node for each variable for each time step.
- Intra-slice arcs: $X_i^T \to X_j^T$
- Inter-slice (temporal) arcs

$$X_i^T \to X_i^{t+1}$$
$$X_i^T \to X_j^{t+1}$$

Fever dynamic Bayesian network

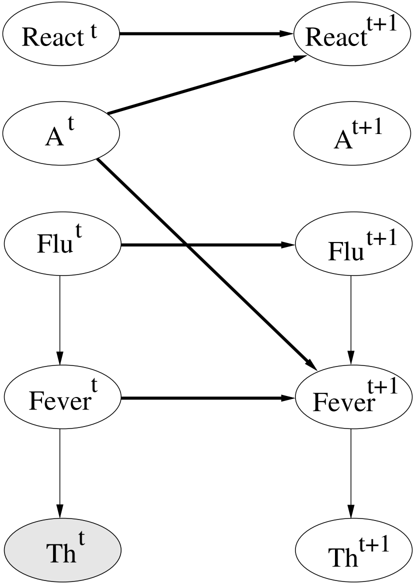

DBN reasoning

- Can calculate distributions for S_{t+1} and further: **probabilistic projection**.
- Reasoning can be done using standard BN updating algorithms
- This type of DBN gets very large, very quickly.
- Usually only keep two time slices of the network.

Dynamic Decision Network

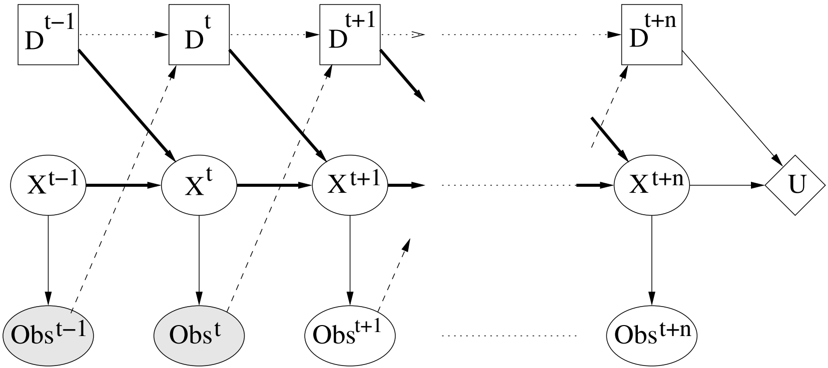

- Similarly, Decision Networks can be extended to include temporal aspects.
- Sequence of decisions taken = Plan.

Fever Dynamic Decision Network

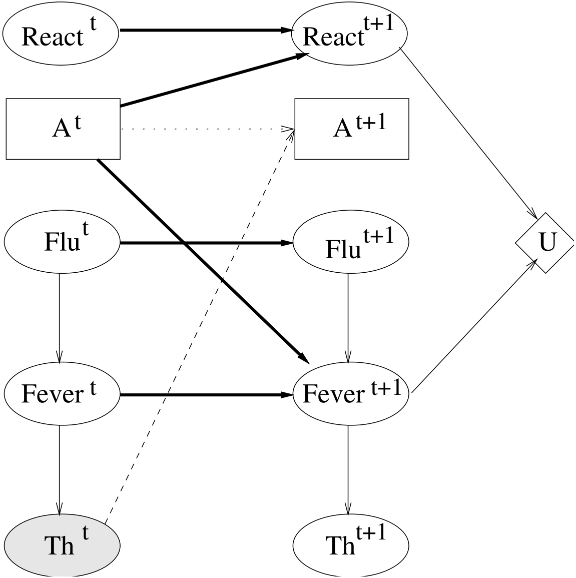

Uses of Bayesian Networks

- Calculating the belief in query variables given values for evidence variables (above).
- Predicting values in dependent variables given values for independent variables.
- Modeling causal interventions.
- Decision making based on probabilities in the network and on the agent's utilities (Influence Diagrams [Howard and Matheson 1981]).
- Deciding which additional evidence should be observed in order to gain useful information.
- Sensitivity analysis to test impact of changes in probabilities or utilities on decisions.

Outline

Summary

- BNs can be extended with decision nodes and utility nodes to support decision making: Decision Networks or Influence Diagrams.
- BNs and decision networks can be extended to allow explicit reasoning about changes over time.

Acknowledgments

Lecture 8 is extracted from
`http://www.csse.monash.edu.au/courseware/cse458/L3-4.pdf`, and composed of materials from [Korb and Nicholson, 2003, Chapter 4] and [Jensen and Nielsen, 2007, Chapter 9] with the instructor's own interpretations. The instructor takes full responsibility of any mistakes in the slides.

References I

Finn V. Jensen and Thomas D. Nielsen.
Bayesian Networks and Decision Graphs.
Springer Publishing Company, Incorporated, 2007.

K. Korb and A. E. Nicholson.
Bayesian Artificial Intelligence.
Chapman & Hall /CRC, 2003.