



## Lecture 15: Uncertainty - 6

Victor Lesser

CMPSCI 683  
Fall 2004

## The value of information

- **Example 1:** You consider buying a program to manage your finances that costs \$100. There is a prior probability of 0.7 that the program is suitable in which case it will have a positive effect on your work worth \$500. There is a probability of 0.3 that the program is not suitable in which case it will have no effect.
- What is the value of knowing whether the program is suitable before buying it?

V. Lesser CS683 F2004

2

## Example 1 Answer

- Expected utility given information  
–  $[0.7*(500-100)+0.3(0)]$
- Expected utility not given information  
–  $[0.7(500-100)+0.3(0-100)]$
- Value of Information  
–  $[0.7*(500-100)+0.3(0)] - [0.7(500-100)+0.3(0-100)] = 280 - 250 = \$30$

V. Lesser CS683 F2004

3

## Value of Perfect Information

- The general case: We assume that exact evidence can be obtained about the value of some random variable  $E_j$ .
- The agent's current knowledge is  $E$ .
- The value of the current best action  $\alpha$  is defined by:

$$EU(\alpha|E) = \max_A \sum_i P(\text{Result}_i(A)|\text{Do}(A),E) U(\text{Result}_i(A))$$

V. Lesser CS683 F2004

4

## VPI cont.

- With the information, the value of the new best action will be:  $EU(\alpha_{E_j} | E, E_j) = \max_A \sum_i P(\text{Result}_i(A) | \text{Do}(A), E, E_j) U(\text{Result}_i(A))$
- But  $E_j$  is a random variable whose value is currently unknown, so we must average over all possible values  $e_{jk}$  using our current belief:  $VPI_E(E_j) = (\sum_k \underline{P(E_j=e_{jk} | E)} EU(\alpha_{e_{jk}} | E, E_j = e_{jk})) - EU(\alpha | E)$

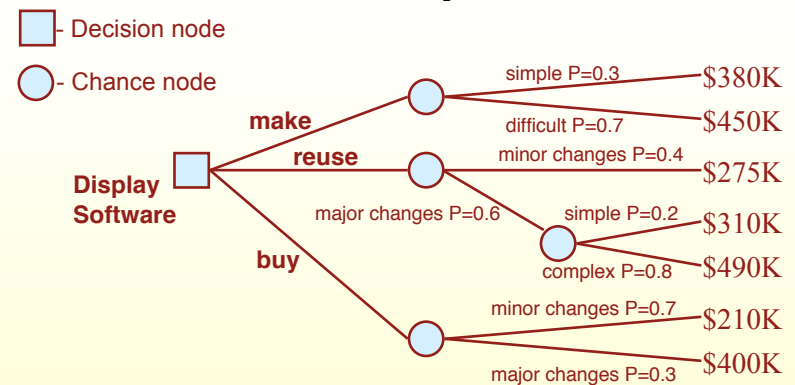
## Outline

- Decision Trees
- Decision Networks
- Markov Decision Processes (MDPs)

## Decision Trees

- A decision tree is an explicit representation of all the possible scenarios from a given state.
- Each path corresponds to decisions made by the agent, actions taken, possible observations, state changes, and a final outcome node.
- Similar to a game played against “nature”

## Example 1: Software Development



- $EU(\text{make}) = 0.3 * \$380K + 0.7 * \$450K = \$429K$ ; **best choice**
- $EU(\text{reuse}) = 0.4 * \$275K + 0.6 * [0.2 * \$310K + 0.8 * \$490K] = \$382.4K$
- $EU(\text{buy}) = 0.7 * \$210K + 0.3 * \$400K = \$267K$

## Example 2: Buying a car


- There are two candidate cars  $C_1$  and  $C_2$ , each can be of good quality (+) or bad quality (-).
- There are two possible tests,  $T_1$  on  $C_1$  (costs \$50) and  $T_2$  on  $C_2$  (costs \$20).
- $C_1$  costs \$1500 (\$500 below market value) but if it is of bad quality repair cost is \$700.
  - 500 gain or 200 lost
- $C_2$  costs \$1150 (\$250 below market value) but if it is of bad quality repair cost is \$150.
  - 250 gain or 100 gain
- Buyer must buy one of the cars and can perform at most one test. -- What other information?

## Example 2: Buying a car cont.

- The chances that the cars are of good quality are 0.70 for  $C_1$  and 0.80 for  $C_2$ .
- Test  $T_1$  will confirm good quality with probability 0.80 and will confirm bad quality with probability 0.65.
- Test  $T_2$  will confirm good quality with probability 0.75 and will confirm bad quality with probability 0.70.

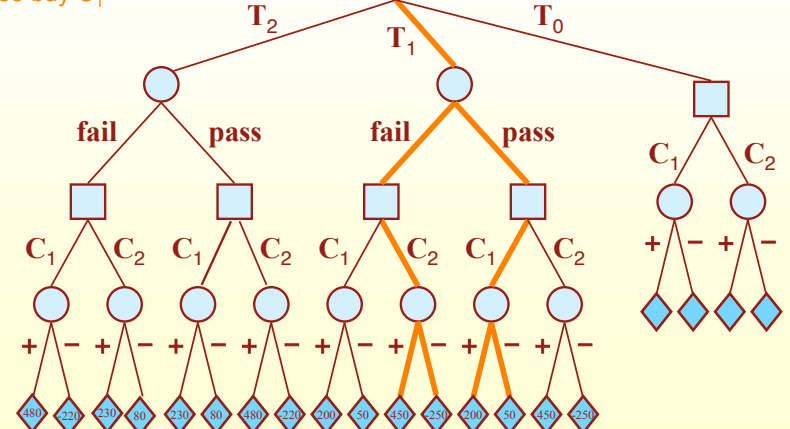
## Example 2: Buying a car cont.

What are the decisions and how can you judge their outcomes?



## Example 2: Buying a car cont.


Do Test  $T_1$ ; If  $T_1$  fails buy  $C_2$  else buy  $C_1$



# Evaluating decision trees

1. Traverse the tree in a depth-first manner:
  - (a) Assign a value to each leaf node based on the outcome
  - (b) Calculate the average utility at each chance node
  - (c) Calculate the maximum utility at each decision node, while marking the maximum branch
2. Trace back the marked branches, from the root node down to find the desired optimal (conditional) plan.

Finding the value of (perfect or imperfect) information in a decision tree.



# Additional Information

Buyer knows car  $c_1$  is good quality

$$70\% P(c_1=\text{good}) = .7$$

Buyer knows car  $c_2$  is good quality

$$80\% P(c_2=\text{good}) = .8$$

Test  $t_1$  check quality of car  $c_1$

$$P(t_1=\text{pass}/c_1=\text{good}) = .8$$

$$P(t_1=\text{pass}/c_1=\text{bad}) = .35$$


Test of  $t_2$  check quality of car  $c_2$

$$P(t_2=\text{pass}/c_2=\text{good}) = .75$$

$$P(t_2=\text{pass}/c_2=\text{bad}) = .3$$


# Details of Example

- **Case 1**
  - $P(c_1=\text{good}/t_2=\text{fail})=p(c_1=\text{good})=.7$
  - Utility =  $2000-1500-20 = 480$
- **Case 2**
  - $P(c_1=\text{bad}/t_2=\text{fail}) = p(c_1=\text{bad}) = 1 - p(c_1=\text{good}) = .3$
  - Utility =  $2000-1500-700-20 = -220$
- **Expected Utility of Chance Node of 1&2**
  - $.7 \times 480 + .3 \times -220 = 270$



# Details of Example cont

- **Case 3**
  - $P(c_2=\text{good}/t_2=\text{fail}) =$
  - $P(t_2=\text{fail}/c_2=\text{good}) P(c_2=\text{good})/P(t_2=\text{fail}) =$
  - $(.25 \times .8 = .2) / P(t_2=\text{fail}) =$
  - Normalize  $.2/.34, .14/.34$  (over  $c_2$  bad)
  - $.59$
  - Utility =  $1400-1150-20 = 230$
- **Case 4**
  - $P(c_2=\text{bad}/t_2=\text{fail}) =$
  - $P(t_2=\text{fail}/c_2=\text{bad}) P(c_2=\text{bad})/P(t_2=\text{fail}) =$
  - $(.7 \times .2 = .14) / P(t_2=\text{fail}) =$
  - $.41$
  - Utility =  $1400-1150-20-150 = 80$
- **Expected Utility of Chance Node of 3&4**
  - $.59 \times 230 + .41 \times 80 = 168.5$






## Details of Example cont

- **What is the decision if**

- Decide to do test t2
- It comes out false
- Do you buy c1 or c2?

- $E(c1/test\ t2=fail)$  = Expected Utility of Chance Node of 1&2 = 270
- $E(c2/test\ t2=fail)$  = Expected Utility of Chance Node of 3&4 = 168.5



## Markov Decision Problems

- **A model of sequential decision-making developed in operations research in the 1950's.**
- **Allows reasoning about actions with uncertain outcomes.**
- **MDPs have been adopted by the AI community as a framework for:**
  - Decision-theoretic planning (e.g., [Dean et al., 1995])
  - Reinforcement learning (e.g., [Barto et al., 1995])

## Markov decision processes

- $S$  - finite set of domain states
- $A$  - finite set of actions
- $P(s'|s,a)$  - state transition function
- $r(s,a)$  - reward function; can get reward at any point
- $S_0$  - initial state
- **The Markov assumption:**

$$P(s_t | s_{t-1}, s_{t-2}, \dots, s_1, a) = P(s_t | s_{t-1}, a)$$

## Example: An Optimal Policy


A policy is a choice of what action to choose at each state

An Optimal Policy is a policy where you are always choosing the action that maximizes the "return"/"utility" of the current state


→	→	→	+1	.812	.868	.912	+1
↑		↑	-1	.762		.660	-1
↑	←	←	←	.705	.655	.611	.388

Actions succeed with probability 0.8 and move at right angles with probability 0.1 (remain in the same position when there is a wall). Actions incur a small cost (0.04).


## Possible Policy Structures



Solution is a **simple path**  
**deterministic**



Solution is an **acyclic graph**  
**Non-deterministic**  
**Based on action**  
**outcomes**



Solution is a **cyclic graph**  
**Allows for infinite**  
**sequence of**  
**action**

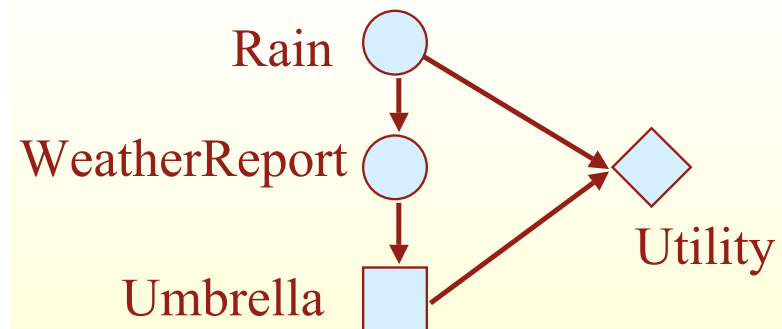
## Decision Networks/Influence Diagrams

- Decision networks or influence diagrams are an extension of belief networks that allow for reasoning about actions and utility.
- The network represents information about the agent's current state, its possible actions, the possible outcome of those actions, and their utility.

## Influence Diagrams

- **Decision trees are not convenient for representing domain knowledge**
  - Requires tremendous amount of storage
    - Multiple decisions nodes -- expands tree
    - Duplication of knowledge along different paths
      - Joint Probability Distribution vs Bayes Net
- **Generate decision tree on the fly from more economical forms of knowledge**
  - Depth-first expansion of tree for computing optimal decision

## Example 3: Taking an Umbrella



Parameters:  $P(\text{Rain})$ ,  $P(\text{WeatherReport}|\text{Rain})$ ,  
 $P(\text{WeatherReport}|\neg\text{Rain})$ ,  $\text{Utility}(\text{Rain}, \text{Umbrella})$

## Nodes in a Decision Network

- **Chance nodes** (ovals) have CPTs (conditional probability tables) that depend on the states of the parent nodes (chance or decision).
- **Decision nodes** (squares) represent options available to the decision maker.
- **Utility nodes** (Diamonds) or value nodes represent the overall utility based on the states of the parent nodes.

## Knowledge in an Influence Diagram

- **Causal knowledge about how events influence each other in the domain**
- **Knowledge about what action sequences are feasible in any given set of circumstances**
  - Lays out possible temporal ordering of decisions
- **Normative knowledge about how desirable the consequences are**

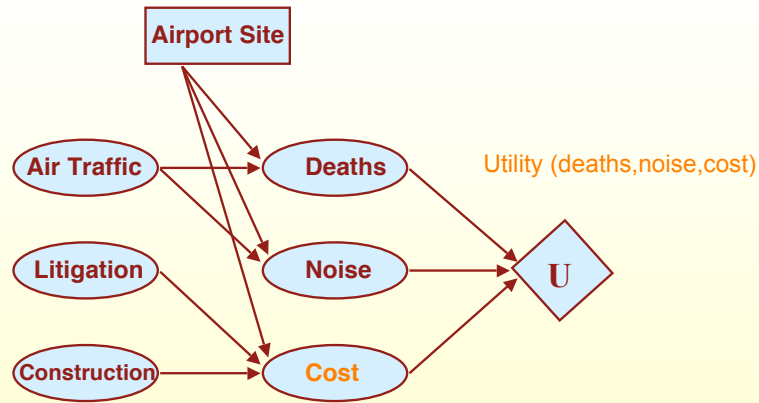
## Topology of decision networks

1. The directed graph has no cycles.
2. The utility nodes have no children.
3. There is a directed path that contains all of the decision nodes.
4. A CPT is attached to each chance node specifying  $P(A|\text{parents}(A))$ .
5. A real valued function over  $\text{parents}(U)$  is attached to each utility node.

## Semantics

- Links into decision nodes are called “information links,” and they indicate that the state of the parent is known prior to the decision.
- The directed path that goes through all the decision nodes defines a temporal sequence of decisions.
- It also partitions the chance variables into sets:  $I_0$  is the vars observed before any decision is made,  $I_1$  is the vars observed after the first and before the second decision, etc.  $I_n$  is the set of unobserved vars.
- The “no-forgetting” assumption is that the decision maker remembers all past observations and decisions. -- **Non Markov Assumption**

## Example 4: Airport Siting Problem



- $P(\text{cost=high}/\text{airportsite=Darien}, \text{airtraffic=low}, \text{litigation=high}, \text{construction=high})$

## Evaluating Decision Networks

1. Set the evidence variables for the current state.
2. For each possible value of the decision node:
  - (a) Set the decision node to that value.
  - (b) Calculate the posterior probabilities for the parent nodes of the utility node.
  - (c) Calculate the expected utility for the action.
3. Return the action with the highest utility.

Similar to Cutset Conditioning of a Multiply Connected Belief Network


## Example 5: Mildew

Two months before the harvest of a wheat field, the farmer observes the state  $Q$  of the crop, and he observes whether it has been attacked by mildew,  $M$ . If there is an attack, he will decide on a treatment with fungicides.

There are five variables:

- $Q$ : fair (f), not too bad (n), average (a), good (g)
- $M$ : no (no), little (l), moderate (m), severe (s)
- $H$ : state of  $Q$  plus  $M$ : rotten (r), bad (b), poor (p)
- $OQ$ : observation of  $Q$ ; imperfect information on  $Q$
- $OM$ : observation of  $M$ ; imperfect information on  $M$

## Mildew decision model



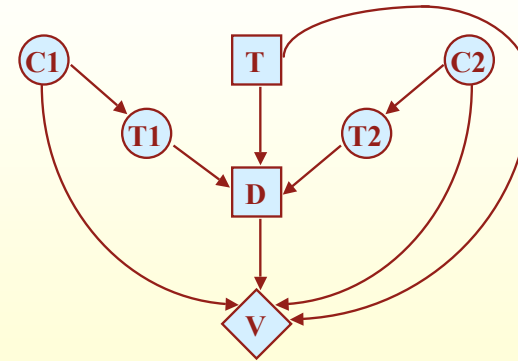


## One action in general

- A single decision node  $D$  may have links to some chance nodes.
- A set of utility functions  $U_1, \dots, U_n$  over domains  $X_1, \dots, X_n$ .
- Goal: find the decision  $d$  that maximizes  $EU(D | e)$ :
- How to solve such problems using a standard Bayesian network package?

$$EU(D | e) = \sum_{X_1} U_1(X_1)P(X_1 | D, e) + \dots + \sum_{X_n} U_n(X_n)P(X_n | D, e)$$

## Multiple decisions -- Policy Generation



Need a more complex evaluation technique since generating a policy

## Options At Decision Node D

- **If T=no test**
  - {Buy 1, Buy 2}
- **If T=do test t1**
  - if t1=pass Buy1 else if t1=fail Buy2
  - if t1=pass Buy2 else if t1=fail Buy1
  - Buy1
  - Buy2
- **If T=do test t2**
  - Same as above

A POLICY IS A SEQUENTIAL SET OF DECISIONS, EACH POTENTIALLY BASED ON THE OUTCOME OF PREVIOUS DECISIONS

## Evaluation by Graph Reduction

Basic idea: (Ross Shachter) Perform a sequence of transformations to the diagram that preserve the optimal policy and its value, until only the UTILITY node remains.

- Similar to ideas of transformation into polytree


Four basic value/utility-preserving reductions:

- Barren node removal
- Chance node removal (marginalization)
- Decision node removal (maximization)
- Arc reversal (Bayes' rule)

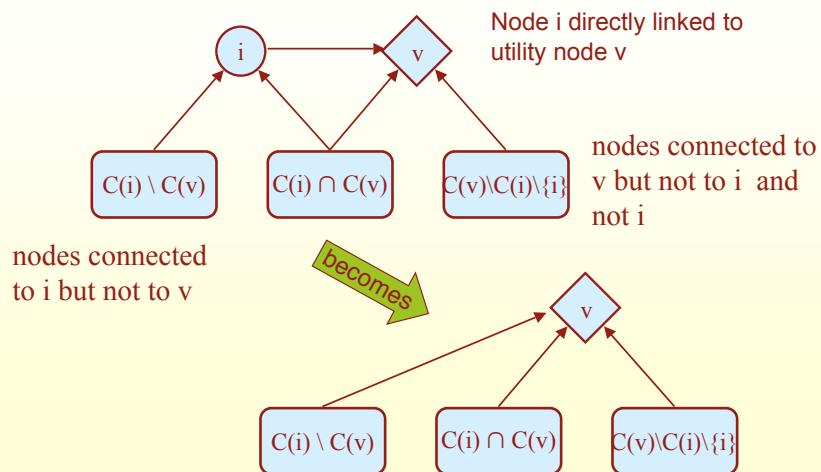
## Barren node reduction

- Let  $X_j$  represent a subset of nodes of interest in an influence diagram.
- Let  $X_k$  represent a subset of evidence nodes.
- We are interested in  $P(f(X_j) | X_k)$
- A node is “barren” if it has no successors and it is not a member of  $X_j$  or  $X_k$ .
- The elimination of barren nodes does not affect the value of  $P(f(X_j) | X_k)$

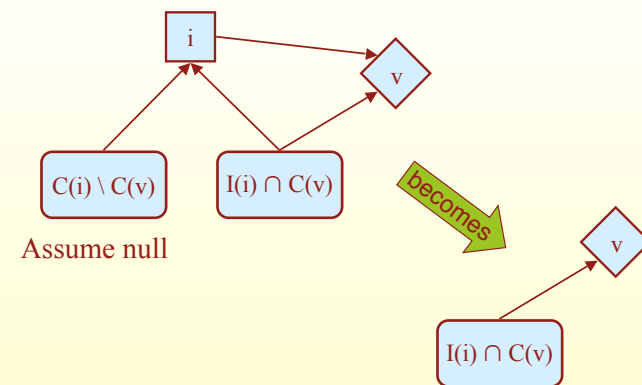
## Barren Node Removal



## Chance Node Removal




## Decision Node Removal



# Arc reversal


- Given an influence diagram containing an arc from  $i$  to  $j$ , but no other directed path from  $i$  to  $j$ , it is possible to transform the diagram to one with an arc from  $j$  to  $i$ . (If  $j$  is deterministic, then it becomes probabilistic.)

## 1 Barren Node Removal




$$\sum_D P(A, B, C, D) = P(A)P(B|A)P(C|B, A) \sum_D P(D|C) \stackrel{1}{=} 1$$

## 1 Removal into Value Node (by Expectation)




$$V(D) = \sum_X V(X, D) * P(X)$$


# Arc Reversal



## Arc Reversal



$$X = Pa(A) \setminus Pa(B) \quad Y = Pa(A) \vee Pa(B) \quad Z = Pa(B) \setminus Pa(A)$$



$$P(A|B, X, Y, Z) = P(B|A, Y, Z) * P(A|X, Y) / P(B|X, Y, Z)$$

$Pa =$  Parents  $Pa(A) \setminus Pa(B)$  parents of A who are not parents of B

## Decision Example

- 1 Reverse X-Y Arc
- 1 Removal of X by Expectation into V
- 1 Removal of D by maximization into V
  - 1 gives you the optimal policy for D given Y

