Introduction to Decision Networks

Alice Gao

Lecture 13

Based on work by K. Leyton-Brown, K. Larson, and P. van Beek

Outline

Learning Goals

By the end of the lecture, you should be able to

- Model a one-off decision problem by constructing a decision network containing nodes, arcs, conditional probability distributions, and a utility function.
- Choose the best action by evaluating a decision network.

Decision Theory

Decision theory = Probability theory + Utility theory

- Decision theory: describes what an agent should do
- Probability theory: describes what an agent should believe on the basis of the evidence
- Utility theory: describes what an agent wants

Decision Networks

Decision networks
$=$ *Bayesian network* $+$ *actions* $+$ *utilities*

A robot that delivers mail

The robot must choose its route to pickup the mail. There is a short route and a long route. The long route is slower, but on the short route the robot might slip and fall. The robot can put on pads. This won't change the probability of an accident, but it will make it less severe if it happens. Unfortunately, the pads add weight and slow the robot down. The robot would like to pick up the mail quickly with little/no damage.

What should the robot do?

Variables

What are the random variables?

What are the decision variables (actions)?

Nodes in a Decision Network

Three kinds of nodes:

- 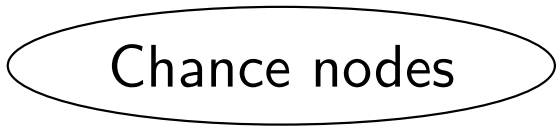

represent random variables (as in Bayesian networks).
- 

represent actions (decision variables).
- 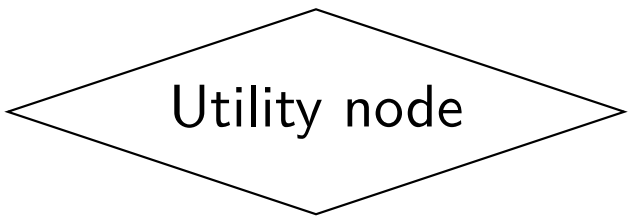

represents agent's utility function on states (happiness in each state).

Robot decision network

Arcs in the Decision Network

How do the random variables and the decision variables relate to one another?

Robot decision network

CQ: The robot's happiness

CQ: Which variables directly influence the robot's happiness?

(A) P only
(B) S only
(C) A only
(D) Two of (A), (B), and (C)
(E) All of (A), (B) and (C)

The robot must choose its route to pickup the mail. There is a short route and a long route. The long route is slower, but on the short route the robot might slip and fall. The robot can put on pads. This won't change the probability of an accident, but it will make it less severe if it happens. Unfortunately, the pads add weight and slow the robot down. The robot would like to pick up the mail quickly with little/no damage.

Robot decision network

CQ The robot's utility function

CQ: When an accident does not happen, which of the following is true?

(A) The robot prefers not wearing pads than wearing pads.

(B) The robot prefers the long route over the short route.

(C) Both (A) and (B) are true.

(D) Neither (A) and (B) is true.

The robot's utility function

	State	$U(w_i)$
$\neg P, \neg S, \neg A$	w_0 slow, no weight	6
$\neg P, \neg S, A$	w_1 impossible	
$\neg P, S, \neg A$	w_2 quick, no weight	10
$\neg P, S, A$	w_3 severe damage	0
$P, \neg S, \neg A$	w_4 slow, extra weight	4
$P, \neg S, A$	w_5 impossible	
$P, S, \neg A$	w_6 quick, extra weight	8
P, S, A	w_7 moderate damage	2

The robot's utility function

How does the robot's utility/happiness depend on the random variables and the decision variables?

The robot's utility function

How does the robot's utility/happiness depend on the random variables and the decision variables?

Robot decision network

Evaluating a decision network

How do we choose an action?

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

Revisiting the Learning Goals

By the end of the lecture, you should be able to

- Model a one-off decision problem by constructing a decision network containing nodes, arcs, conditional probability distributions, and a utility function.
- Choose the best action by evaluating a decision network.