CS 188: Artificial Intelligence
Spring 2009

Lecture 20: Decision Networks
4/2/2009

John DeNero – UC Berkeley
Slides adapted from Dan Klein

Announcements

- Written 3 released tonight, due April 14

Decision Networks

- MEU: choose the action which maximizes the expected utility given the evidence
- Can directly operationalize this with decision networks
  - Bayes nets with nodes for utility and actions
  - Lets us calculate the expected utility for each action
- New node types:
  - Chance nodes (just like BNs)
  - Actions (rectangles, must be parents, act as observed evidence)
  - Utility node (diamond, depends on action and chance nodes)

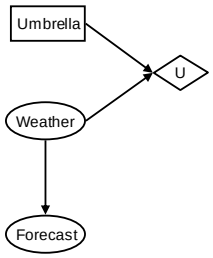

Decision Networks

- Action selection:
  - Instantiate all evidence
  - Set action node(s) each possible way
  - Calculate posterior for all parents of utility node, given the evidence
  - Calculate expected utility for each action
  - Choose maximizing action

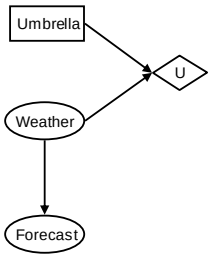

Example: Decision Networks

Umbrella = leave

$$\mathrm{EU}(\text{leave}) = \sum_w P(w)U(\text{leave}, w)$$

$$= 0.7 \cdot 100 + 0.3 \cdot 0 = 70$$

Umbrella = take

$$\mathrm{EU}(\text{take}) = \sum_w P(w)U(\text{take}, w)$$

$$= 0.7 \cdot 20 + 0.3 \cdot 70 = 35$$

Optimal decision = leave

$$\mathrm{MEU}(\emptyset) = \max_a \mathrm{EU}(a) = 70$$

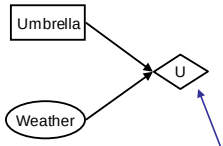

W	P(W)
sun	0.7
rain	0.3

A	W	U(A,W)
leave	sun	100
leave	rain	0
take	sun	20
take	rain	70

Evidence in Decision Networks

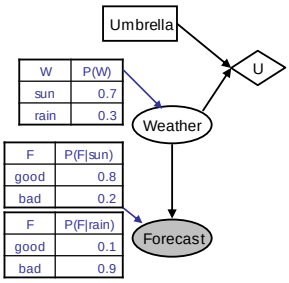

W	P(W)
sun	0.7
rain	0.3

F	P(F|sun)
good	0.8
bad	0.2

F	P(F|rain)
good	0.1
bad	0.9

- Find P(W|F=bad)
  - Select for evidence

W	P(W)
sun	0.7
rain	0.3

$P(W)$

W	P(F=bad|W)
sun	0.2
rain	0.9

$P(bad|W)$

  - First we join P(W) and P(bad|W)
  - Then we normalize

W	P(W,F=bad)
sun	0.14
rain	0.27

$P(W, bad)$

W	P(W | F=bad)
sun	0.34
rain	0.66

$P(W|F = \text{bad})$

Example: Decision Networks

Umbrella = leave

$$\mathrm{EU}(\mathrm{leave}|\mathrm{bad}) = \sum_w P(w|\mathrm{bad})U(\mathrm{bad}, w)$$

$$= 0.34 \cdot 100 + 0.66 \cdot 0 = 34$$

Umbrella = take

$$\mathrm{EU}(\mathrm{take}|\mathrm{bad}) = \sum_w P(w|\mathrm{bad})U(\mathrm{take}, w)$$

$$= 0.34 \cdot 20 + 0.66 \cdot 70 = 53$$

Optimal decision = take

$$\mathrm{MEU}(F = \mathrm{bad}) = \max_a \mathrm{EU}(a|\mathrm{bad}) = 53$$

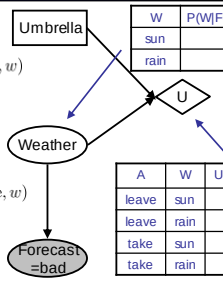

W	P(W|F=bad)
sun	0.34
rain	0.66

A	W	U(A,W)
leave	sun	100
leave	rain	0
take	sun	20
take	rain	70

[Demo]

Conditioning on Action Nodes

- An action node can be a parent of a chance node
- Chance node conditions on the outcome of the action
- Action nodes are like observed variables in a Bayes' net, except we max over their values

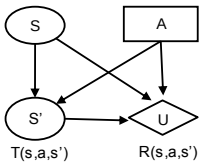

Value of Information

- Idea: compute value of acquiring each possible piece of evidence
  - Can be done directly from decision network
- Example: buying oil drilling rights
  - Two blocks A and B, exactly one has oil, worth k
  - Prior probabilities 0.5 each, & mutually exclusive
  - Drilling in either A or B has MEU = k/2
  - Fair price of drilling rights: k/2
- Question: what's the value of information
  - Value of knowing which of A or B has oil
  - Value is expected gain in MEU from new info
  - Survey may say "oil in a" or "oil in b," prob 0.5 each
  - If we know OilLoc, MEU is k (either way)
  - Gain in MEU from knowing OilLoc?
  - VPI(OilLoc) = k/2
  - Fair price of information: k/2

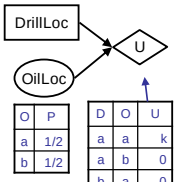

O	P
a	1/2
b	1/2

D	O	U
a	a	k
a	b	0
b	a	0
b	b	k

Value of Perfect Information

- Current evidence E=e, utility depends on S=s

$$\mathrm{MEU}(e) = \max_a \sum_s P(s|e)\, U(s,a)$$

- Potential new evidence E': suppose we knew E' = e'

$$\mathrm{MEU}(e, e') = \max_a \sum_s P(s|e, e')\, U(s,a)$$

- BUT E' is a random variable whose value is currently unknown, so:
  - Must compute expected gain over all possible values

$$\mathrm{VPI}_e(E') = \sum_{e'} P(e'|e) \left(\mathrm{MEU}(e, e') - \mathrm{MEU}(e)\right)$$

- (VPI = value of perfect information)

VPI Example: Weather

MEU with no evidence

$$\mathrm{MEU}(\emptyset) = \max_a \mathrm{EU}(a) = 70$$

MEU if forecast is bad

$$\mathrm{MEU}(F = \mathrm{bad}) = \max_a \mathrm{EU}(a|\mathrm{bad}) = 53$$

MEU if forecast is good

$$\mathrm{MEU}(F = \mathrm{good}) = \max_a \mathrm{EU}(a|\mathrm{good}) = 95$$

Forecast distribution

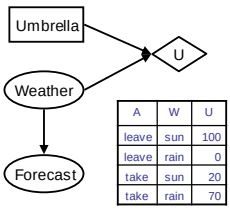

A	W	U
leave	sun	100
leave	rain	0
take	sun	20
take	rain	70

F	P(F)
good	0.59
bad	0.41

$$0.59 \cdot (95 - 70) + 0.41 \cdot (53 - 70)$$

$$0.59 \cdot (+25) + 0.41 \cdot (-17) = \mathbf{7.8}$$

$$\mathrm{VPI}_e(E') = \sum_{e'} P(e'|e) \left(\mathrm{MEU}(e, e') - \mathrm{MEU}(e)\right)$$

VPI Example: Ghostbusters

- Reminder: ghost his hidden, sensors are noisy
- T: Top square is red
  B: Bottom square is red
  G: Ghost is in the top
- Sensor model:

  P(t | g) = 0.8
  P(t | ¬g) = 0.4
  P(b | g) = 0.4
  P(b | ¬g) = 0.8

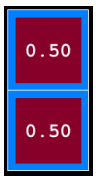

Joint Distribution

T	B	G	P(T,B,G)
t	b	g	0.16
t	b	¬g	0.16
t	¬b	g	0.24
t	¬b	¬g	0.04
¬t	b	g	0.04
¬t	b	¬g	0.24
¬t	¬b	g	0.06
¬t	¬b	¬g	0.06

[Demo]

VPI Example: Ghostbusters

Utility of bust is 2, no bust is 0

- Q1: What's the value of knowing T if I know nothing?
- Q1': $E_{P(T)}$[MEU(t) – MEU()]
- Q2: What's the value of knowing B if I already know that T is true (red)?
- Q2': $E_{P(B|t)}$[MEU(t,b) – MEU(t)]
- How low can the value of information ever be?

Joint Distribution

T	B	G	P(T,B,G)
t	b	g	0.16
t	b	¬g	0.16
t	¬b	g	0.24
t	¬b	¬g	0.04
¬t	b	g	0.04
¬t	b	¬g	0.24
¬t	¬b	g	0.06
¬t	¬b	¬g	0.06

[Demo]

VPI Properties

- Nonnegative in expectation

$$\forall E', e : \mathrm{VPI}_e(E') \geq 0$$

- Nonadditive ---consider, e.g., obtaining E_j twice

$$\mathrm{VPI}_e(E_j, E_k) \neq \mathrm{VPI}_e(E_j) + \mathrm{VPI}_e(E_k)$$

- Order-independent

$$\begin{aligned}\mathrm{VPI}_e(E_j, E_k) &= \mathrm{VPI}_e(E_j) + \mathrm{VPI}_{e,E_j}(E_k) \\ &= \mathrm{VPI}_e(E_k) + \mathrm{VPI}_{e,E_k}(E_j)\end{aligned}$$

14

Quick VPI Questions

- The soup of the day is either clam chowder or split pea, but you wouldn't order either one. What's the value of knowing which it is?
- If you have $10 to bet and odds are 3 to 1 that Berkeley will beat Stanford, what's the value of knowing the outcome in advance, assuming you can make a fair bet for either Cal or Stanford?
- What if you are morally obligated not to bet against Cal, but you can refrain from betting?

15