Decisions and Value of Information

Hal Daumé III

Computer Science
University of Maryland

me@hal3.name

CS 421: Introduction to Artificial Intelligence

17 Apr 2012

Many slides courtesy of
Dan Klein, Stuart Russell,
or Andrew Moore

Announcements

- Today:
  - Finish inference in “simple” networks
  - How to make decisions based on probabilistic inference

- Coming soon!
  - Reasoning over time

Recap: Inference Example

- Find P(W|F=bad)
  - Restrict all factors

W	P(W)
sun	0.7
rain	0.3

$P(W)$

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

$P(bad|W)$

  - No hidden vars to eliminate (this time!)
  - Just join and normalize

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

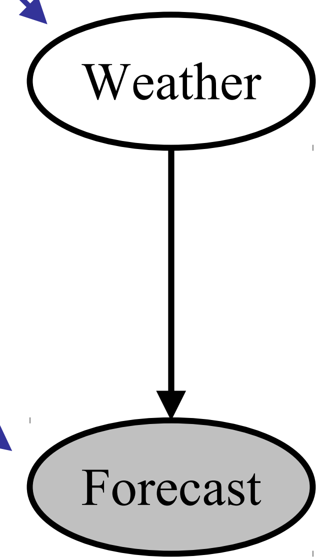

W	P(W,F=bad)
sun	0.14
rain	0.27

$P(W, bad) = P(W) \times P(bad|W)$

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

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

Decision Networks

- MEU: choose the action which maximizes the expected utility given the evidence

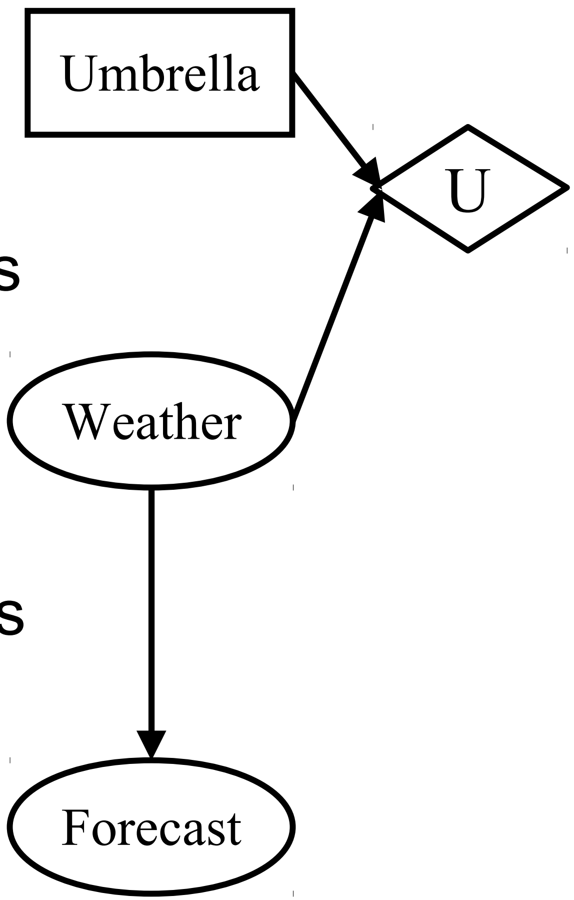

- Can directly operationalize this with decision diagrams
  - Bayes nets with nodes for utility and actions
  - Lets us calculate the expected utility for each action
- New node types:
  - Chance nodes (just like BNs)
  - Actions (rectangles, must be parents, act as observed evidence)
  - Utilities (depend on action and chance nodes)

Decision Networks

- Action selection:
  - Instantiate all evidence
  - Calculate posterior over parents of utility node
  - Set action node each possible way
  - Calculate expected utility for each action
  - Choose maximizing action

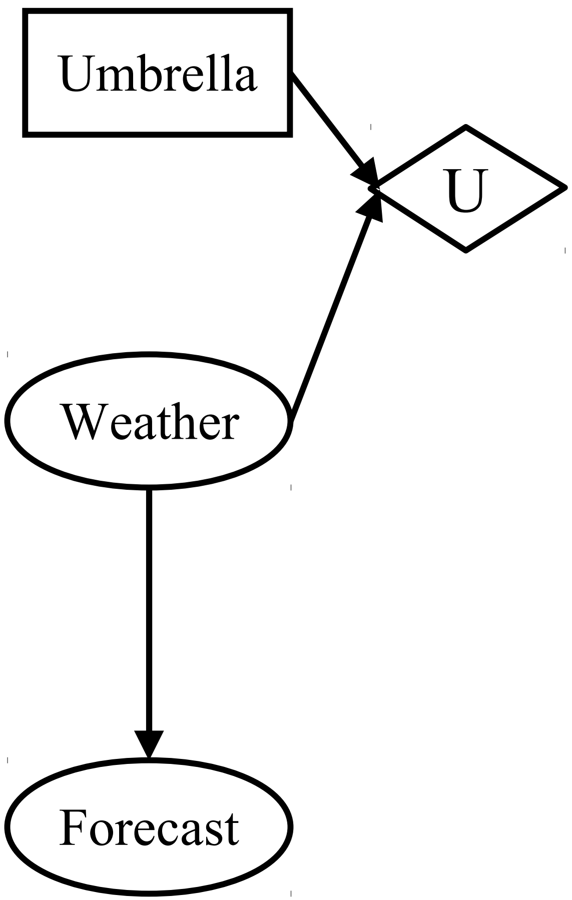

Example: Decision Networks

Umbrella = leave

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

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

Umbrella = take

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

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

Optimal = leave

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

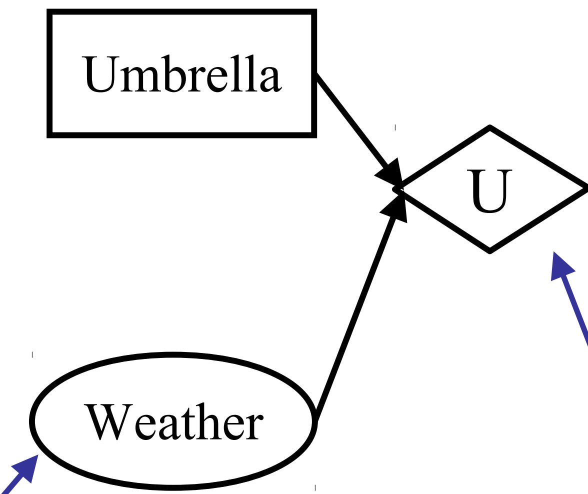

W	P(W)
sun	0.7
rain	0.3

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

Example: Decision Networks

Umbrella = leave

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

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

Umbrella = take

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

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

Optimal = take

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

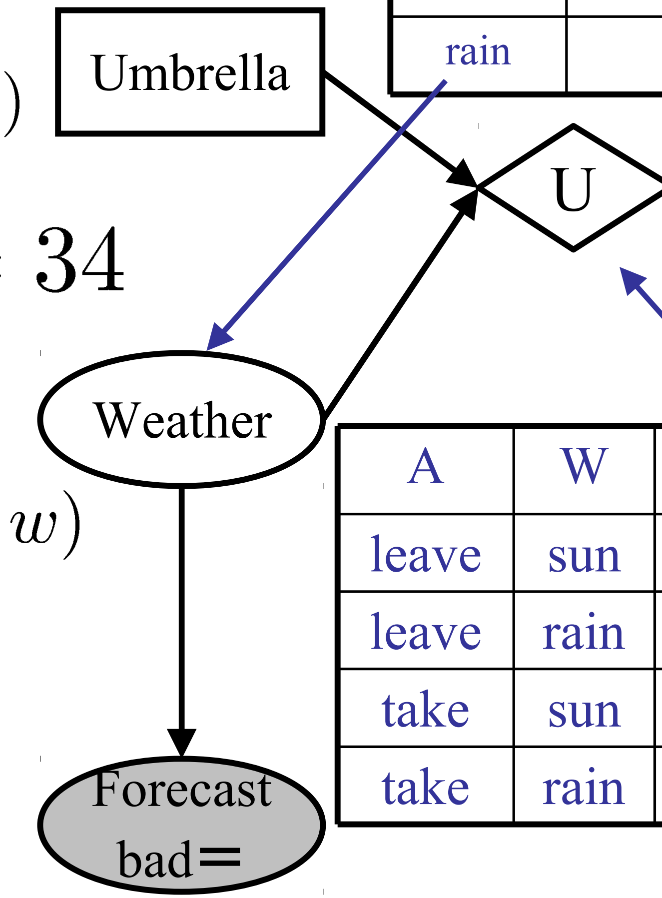

W	P(W|F=bad)
sun	0.34
rain	0.66

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

Value of Information

- Idea: compute value of acquiring each piece of evidence
  - Can be done directly from decision network

- Example: buying oil drilling rights
  - Two blocks A and B, !one has oil, worth k
  - Prior probabilities 0.5 each
  - Current price of each block is k/2
  - MEU = 0 (either action is a maximizer)

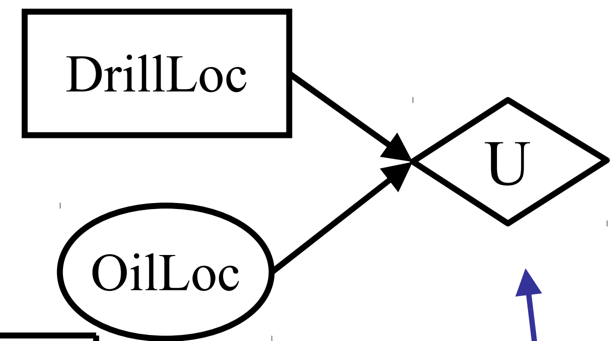

O	P
a	1/2
b	1/2

D	O	U
a	a	k/2
a	b	-k/2
b	a	-k/2
b	b	k/2

- Solution: compute value of information
  = expected gain in MEU from observing new information
- Probe gives accurate survey of A.
- Fair price?
  - Survey may say “oil in a” or “oil in b,” prob 0.5 each
  - If we know O, MEU is k/2 (either way)
  - Gain in MEU?
  - VPI(O) = k/2
  - Fair price: k/2

Value of Information

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

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

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

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

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

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

- (VPI = value of perfect information)

VPI Example

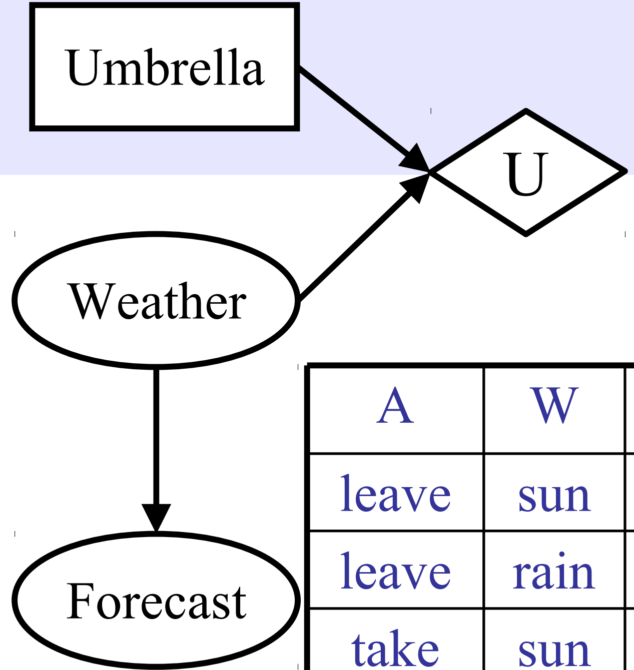

A	W	U(A,W)
leave	sun	100
leave	rain	0
take	sun	20
take	rain	70

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

F	P(F)
good	0.59
bad	0.41

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

$$0.59 \cdot (+25) + 0.41 \cdot (-17) = +22$$

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

VPI Properties

- Nonnegative in expectation

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

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

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

- Order-independent

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

VPI Scenarios

- Imagine actions 1 and 2, for which $U_1 > U_2$
- How much will information about E_j be worth?

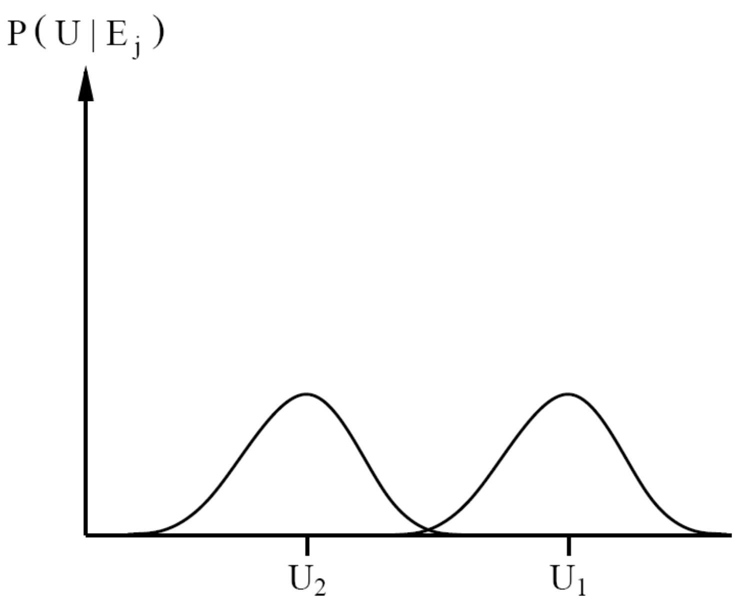

Little – we're sure action 1 is better.

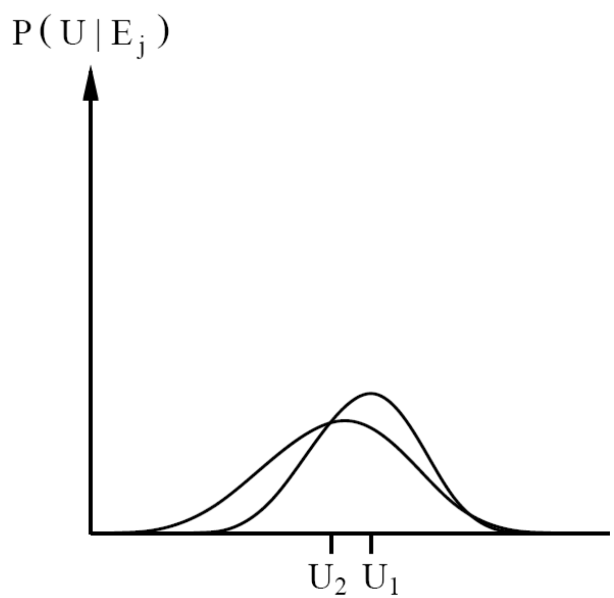

A lot – either could be much better

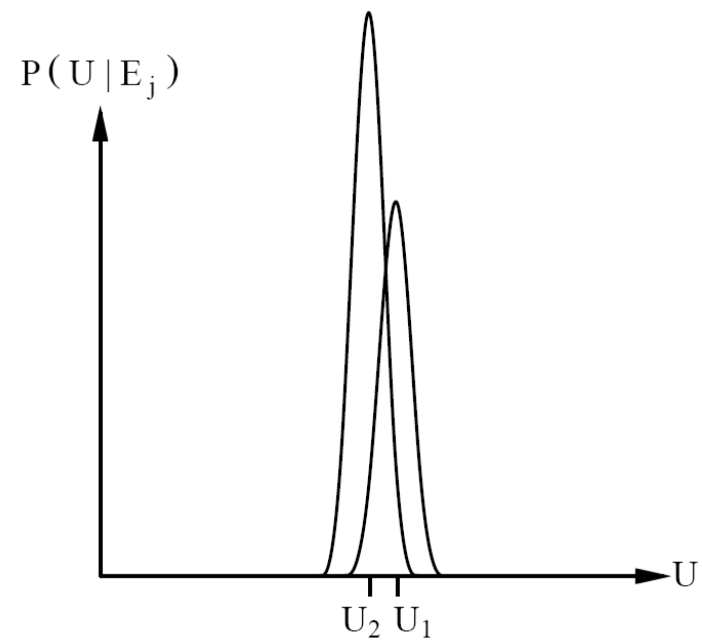

Little – info likely to change our action but not our utility