

CS 188: Artificial Intelligence Spring 2010

Lecture 19: Decision Diagrams 4/1/2010

Pieter Abbeel – UC Berkeley
Many slides over this course adapted from Dan Klein, Stuart Russell,
Andrew Moore


Announcements

- Mid-Semester Evaluations
 - Link is in your email
- Assignments
 - W5 due tonight
 - W6 out tonight

2

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, cannot have parents, act as observed evidence)
 - Utility node (diamond, depends on action and chance nodes)



4

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



5

Value of Information

- Idea: compute value of acquiring evidence
 - Can be done directly from decision network
- Example: buying oil drilling rights
 - Two blocks A and B, exactly one has oil, worth k
 - You can drill in one location
 - Prior probabilities 0.5 each, & mutually exclusive
 - Drilling in either A or B has MEU = 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



12

Value of Information


- Assume we have evidence E=e. Value if we act now:

$$MEU(e) = \max_a \sum_s P(s|e) U(s, a)$$
- Assume we see that E' = e'. Value if we act then:

$$MEU(e, e') = \max_a \sum_s P(s|e, e') U(s, a)$$
- BUT E' is a random variable whose value is unknown, so we don't know what e' will be
- Expected value if E' is revealed and then we act:

$$MEU(e, E') = \sum_{e'} P(e'|e) MEU(e, e')$$
- Value of information: how much MEU goes up by revealing E' first:

$$VPI(E'|e) = MEU(e, E') - MEU(e)$$



VPI Example: Weather

MEU with no evidence

- MEU(\emptyset) = $\max_a EU(a) = 70$

MEU if forecast is bad

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

MEU if forecast is good

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

Forecast distribution

F	P(F)
good	0.59
bad	0.41

$$0.59 \cdot (95) + 0.41 \cdot (53) - 70 = 77.8 - 70 = 7.8$$

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

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

15

VPI Properties

- Nonnegative
 - $\forall E', e : VPI(E'|e) \geq 0$
- Nonadditive --- consider, e.g., obtaining E_i twice
 - $VPI(E_j, E_k|e) \neq VPI(E_j|e) + VPI(E_k|e)$
- Order-independent
 - $VPI(E_j, E_k|e) = VPI(E_j|e) + VPI(E_k|e, E_j)$
 - $= VPI(E_k|e) + VPI(E_j|e, E_k)$

16

Quick VPI Questions

Known answer: +10

- 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?
- There are two kinds of plastic forks at a picnic. It must be that one is slightly better. What's the value of knowing which?
 - But 0: $0.75 \cdot 10 + 0.15 \cdot (-10) = 5$*
 - But 5: $0.75 \cdot 10 + 0.15 \cdot (-10) = 5$*
- You have \$10 to bet double-or-nothing and there is a 75% chance that Berkeley will beat Stanford. What's the value of knowing the outcome in advance?

You must bet on Cal, either way. What's the value now?

Reasoning over Time

- Often, we want to reason about a sequence of observations
 - Speech recognition
 - Robot localization
 - User attention
 - Medical monitoring
- Need to introduce time into our models
- Basic approach: hidden Markov models (HMMs)
- More general: dynamic Bayes' nets

21

Markov Models

- A Markov model is a chain-structured BN
 - Each node is identically distributed (stationarity)
 - Value of X at a given time is called the state
 - As a BN:
 - $X_1 \rightarrow X_2 \rightarrow X_3 \rightarrow X_4 \rightarrow \dots$

$P(X_1) \quad P(X|X_{-1})$

- Parameters: called transition probabilities or dynamics, specify how the state evolves over time (also, initial probs)

Example: Markov Chain

- Weather:
 - States: $X = \{\text{rain}, \text{sun}\}$
 - Transitions:
 - Initial distribution: 1.0 sun
 - What's the probability distribution after one step?

$P(X_2 = \text{sun}) = P(X_2 = \text{sun}|X_1 = \text{sun})P(X_1 = \text{sun}) + P(X_2 = \text{sun}|X_1 = \text{rain})P(X_1 = \text{rain})$

$0.9 \cdot 1.0 + 0.1 \cdot 0.0 = 0.9$

25

Mini-Forward Algorithm

- Question: What's $P(X)$ on some day t ?
 - An instance of variable elimination!

$$\rightarrow P(x_t) = \sum_{x_{t-1}} P(x_t|x_{t-1})P(x_{t-1})$$

$$\rightarrow P(x_1) = \text{known}$$

Forward simulation

Example

- From initial observation of sun

$\begin{pmatrix} 1.0 \\ 0.0 \end{pmatrix}$	$\begin{pmatrix} 0.9 \\ 0.1 \end{pmatrix}$	$\begin{pmatrix} 0.82 \\ 0.18 \end{pmatrix}$	\rightarrow	$\begin{pmatrix} 0.5 \\ 0.5 \end{pmatrix}$
$P(X_1)$	$P(X_2)$	$P(X_3)$		$P(X_\infty)$
- From initial observation of rain

$\begin{pmatrix} 0.0 \\ 1.0 \end{pmatrix}$	$\begin{pmatrix} 0.1 \\ 0.9 \end{pmatrix}$	$\begin{pmatrix} 0.18 \\ 0.82 \end{pmatrix}$	\rightarrow	$\begin{pmatrix} 0.5 \\ 0.5 \end{pmatrix}$
$P(X_1)$	$P(X_2)$	$P(X_3)$		$P(X_\infty)$

Stationary Distributions

- If we simulate the chain long enough:
 - What happens?
 - Uncertainty accumulates
 - Eventually, we have no idea what the state is!
- Stationary distributions:
 - For most chains, the distribution we end up in is independent of the initial distribution
 - Called the **stationary distribution** of the chain
 - Usually, can only predict a short time out

Web Link Analysis

- PageRank** over a web graph
 - Each web page is a state
 - Initial distribution: uniform over pages
 - Transitions:
 - With prob. c , uniform jump to a random page (dotted lines, not all shown)
 - With prob. $1-c$, follow a random outlink (solid lines)
- Stationary distribution
 - Will spend more time on highly reachable pages
 - E.g. many ways to get to the Acrobat Reader download page
 - Somewhat robust to link spam
 - Google 1.0 returned the set of pages containing all your keywords in decreasing rank, now all search engines use link analysis along with many other factors (rank actually getting less important over time)

Hidden Markov Models

- Markov chains not so useful for most agents
 - Eventually you don't know anything anymore
 - Need observations to update your beliefs
- Hidden Markov models (HMMs)
 - Underlying Markov chain over states S
 - You observe outputs (effects) at each time step
 - As a Bayes' net:

Example


R_{t-1}	$P(R_t)$
r	0.7
f	0.3

R_t	$P(U_t)$
r	0.9
f	0.2

- An HMM is defined by:
 - Initial distribution: $P(X_1)$
 - Transitions: $P(X|X_{-1})$
 - Emissions: $P(E|X)$

Ghostbusters HMM

- $P(X_1)$ = uniform
- ▪ $P(X|X')$ = usually move clockwise, but sometimes move in a random direction or stay in place
- ▪ $P(R_{i,j}|X)$ = same sensor model as before: red means close, green means far away.



1/9	1/9	1/9
1/9	1/9	1/9
1/9	1/9	1/9

$P(X_1)$

1/6	1/6	1/2
0	1/6	0
0	0	0

$P(X|X'=\langle 1,2 \rangle)$