

CS 188: Artificial Intelligence Spring 2010

Lecture 15: Bayes' Nets II – Independence 3/9/2010

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

Announcements

- Current readings
 - Require login
- Assignments
 - W4 due Thursday
- Midterm
 - 3/18, 6-9pm, 0010 Evans --- no lecture on 3/18
 - We will be posting practice midterms
 - One page note sheet, non-programmable calculators
 - Topics go through Thursday, not next Tuesday

2

Outline

- Thus far: Probability
- Today: Bayes nets
 - Semantics
 - (Conditional) Independence

3

Probability recap

- Conditional probability $P(x|y) = \frac{P(x,y)}{P(y)}$
- Product rule $P(x,y) = P(x|y)P(y)$
- Chain rule $P(X_1, X_2, \dots, X_n) = P(X_1)P(X_2|X_1)P(X_3|X_1, X_2) \dots$
- X, Y independent iff: $\forall x, y: P(x,y) = P(x)P(y)$
- X and Y are conditionally independent given Z iff:
 $\forall x, y, z: P(x,y|z) = P(x|z)P(y|z)$ $X \perp\!\!\!\perp Y | Z$

4


Bayes' Nets: Big Picture

- Two problems with using full joint distribution tables as our probabilistic models:
 - Unless there are only a few variables, the joint is WAY too big to represent explicitly
 - Hard to learn (estimate) anything empirically about more than a few variables at a time
- Bayes' nets: a technique for describing complex joint distributions (models) using simple, local distributions (conditional probabilities)
 - More properly called graphical models
 - We describe how variables locally interact
 - Local interactions chain together to give global, indirect interactions
 - For about 10 min, we'll be vague about how these interactions are specified

5

Graphical Model Notation

- Nodes: variables (with domains)
 - Can be assigned (observed) or unassigned (unobserved)
- Arcs: interactions
 - Similar to CSP constraints
 - Indicate "direct influence" between variables
 - Formally: encode conditional independence (more later)
- For now: imagine that arrows mean direct causation (in general, they don't!)



8

Example: Coin Flips

- N independent coin flips

X_1 X_2 ... X_n

- No interactions between variables:
absolute independence

9

Example: Traffic

- Variables:
 - R: It rains
 - T: There is traffic
- Model 1: independence
- Model 2: rain causes traffic
- Why is an agent using model 2 better?

10

Example: Traffic II

- Let's build a causal graphical model
- Variables
 - T: Traffic
 - R: It rains
 - L: Low pressure
 - D: Roof drips
 - B: Ballgame
 - C: Cavity

11

Example: Alarm Network

- Variables
 - B: Burglary
 - A: Alarm goes off
 - M: Mary calls
 - J: John calls
 - E: Earthquake!

12

Bayes' Net Semantics

- Let's formalize the semantics of a Bayes' net
- A set of nodes, one per variable X
- A directed, acyclic graph
- A conditional distribution for each node
 - A collection of distributions over X, one for each combination of parents' values
- CPT: conditional probability table
- Description of a noisy "causal" process

$P(X|A_1 \dots A_n)$

13

Probabilities in BNs

- Bayes' nets **implicitly** encode joint distributions
 - As a product of local conditional distributions
 - To see what probability a BN gives to a full assignment, multiply all the relevant conditionals together:

$$P(x_1, x_2, \dots, x_n) = \prod_{i=1}^n P(x_i | \text{parents}(X_i))$$


Example: $P(+cavity, +catch, -toothache)$

Assumption: $P(x_i | x_1, \dots, x_{i-1}) = P(x_i | \text{parents}(x_i))$

- This lets us reconstruct any entry of the full joint
- Not every BN can represent every joint distribution
 - The topology enforces certain conditional independencies

14

Example: Coin Flips




$$P(h, h, t, h) = P(X_1=h) \cdot P(X_2=h) \cdot P(X_3=t) \cdot P(X_4=h)$$

$$= \frac{1}{2} \cdot \frac{1}{2} \cdot \frac{1}{2} \cdot \frac{1}{2}$$

Only distributions whose variables are absolutely independent can be represented by a Bayes' net with no arcs.

15

Example: Traffic

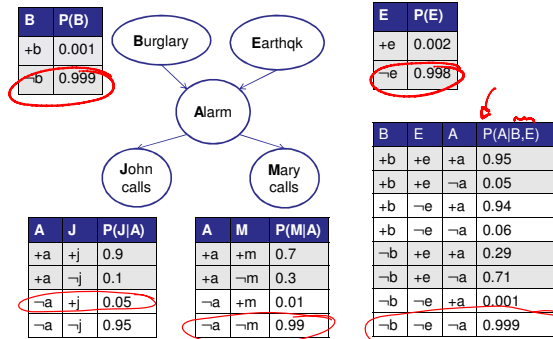


$$P(+r, -t) = P(+r) \cdot P(-t|+r)$$

$$= \frac{1}{4} \cdot \frac{1}{4} = \frac{1}{16}$$

16

Example: Alarm Network



$P(J A)$		$P(M A)$	
+a	+j	0.9	
+a	-j	0.1	
-a	+j	0.05	
-a	-j	0.95	

$P(J A)$		$P(M A)$	
+a	+m	0.7	
+a	-m	0.3	
-a	+m	0.01	
-a	-m	0.99	

$P(J A)$		$P(M A)$	
+a	+j	0.9	
+a	-j	0.1	
-a	+j	0.05	
-a	-j	0.95	

Size of a Bayes' Net

- How big is a joint distribution over N Boolean variables? 2^N
- How big is an N -node net if nodes have up to k parents? $O(N * 2^{k+1})$
- Both give you the power to calculate $P(X_1, X_2, \dots, X_n)$
- BNs: Huge space savings!
- Also easier to elicit local CPTs
- Also turns out to be faster to answer queries (coming)

18

Bayes' Nets

- So far: how a Bayes' net encodes a joint distribution
- Next: how to answer queries about that distribution
 - Key idea: conditional independence
- After that: how to answer numerical queries (inference) more efficiently than by first constructing the joint distribution

19

Conditional Independence


- Reminder: independence
 - X and Y are independent if $\forall x, y, P(x, y) = P(x)P(y) \implies X \perp\!\!\!\perp Y$
 - X and Y are conditionally independent given Z $\forall x, y, z, P(x, y|z) = P(x|z)P(y|z) \implies X \perp\!\!\!\perp Y | Z$
 - (Conditional) independence is a property of a distribution $P(x_i | x_1, \dots, x_{i-1}) = P(x_i | Pa(x_i))$

23

Example: Independence

- For this graph, you can fiddle with θ (the CPTs) all you want, but you won't be able to represent any distribution in which the flips are dependent!


	X_1		X_2	
	$P(X_1)$		$P(X_2)$	
	h	0.5	h	0.5
	t	0.5	t	0.5



24

Topology Limits Distributions


- Given some graph topology G , only certain joint distributions can be encoded
- The graph structure guarantees certain (conditional) independences
- (There might be more independence)
- Adding arcs increases the set of distributions, but has several costs
- Full conditioning can encode any distribution



25

Independence in a BN

- Important question about a BN:
 - Are two nodes independent given certain evidence?
 - If yes, can prove using algebra (tedious in general)
 - If no, can prove with a counter example
 - Example:



- Question: are X and Z necessarily independent?
 - Answer: no. Example: low pressure causes rain, which causes traffic.
 - X can influence Z, Z can influence X (via Y)
 - Addendum: they *could* be independent: how?

Causal Chains

- This configuration is a "causal chain"



X: Low pressure
Y: Rain
Z: Traffic

$$P(x, y, z) = P(x)P(y|x)P(z|y)$$

- Is X independent of Z given Y?

$$P(x, z|y) = \frac{P(x, y, z)}{P(y)} = \frac{P(x)P(y|x)P(z|y)}{P(y)}$$

$$P(x|y, z) = \frac{P(x, z|y)}{P(z|y)} = \frac{P(x)P(y|x)P(z|y)}{P(y)P(z|y)} = \frac{P(x)P(y|x)}{P(y)}$$

- Evidence along the chain "blocks" the influence. Yes!

27

Common Cause

- Another basic configuration: two effects of the same cause

- Are X and Z independent? **NO**
- Are X and Z independent given Y? **Yes!**

$$P(z|x, y) = \frac{P(x, y, z)}{P(x, y)} = \frac{P(y)P(x|y)P(z|y)}{P(y)P(x|y)} = P(z|y)$$


$$\stackrel{?}{=} P(z|y)$$

$$P(x, y, z) = P(y)P(x|y)P(z|y)$$

- Observing the cause blocks influence between effects.

28


$P(x=y)=5$
 $P(x|y)=1$ fact
 $P(z|y)=1$ fact



Common Effect

- Last configuration: two causes of one effect (v-structures)

- Are X and Z independent?
 - Yes: the ballgame and the rain cause traffic, but they are not correlated
 - Still need to prove they must be (try it!)
- Are X and Z independent given Y?
 - No: seeing traffic puts the rain and the ballgame in competition as explanation?
- This is backwards from the other cases
 - Observing an effect *activates* influence between possible causes.




X: Raining
Z: Ballgame
Y: Traffic

29

Causal chain: $0 \rightarrow 0 \rightarrow 0$, common cause $0 \rightarrow 0$

The General Case

- Any complex example can be analyzed using these three canonical cases 
- General question: in a given BN, are two variables independent (given evidence)?
- Solution: analyze the graph

30