

# Bayes Nets II: Independence Day



CS 5300 / CS 6300  
Artificial Intelligence  
Spring 2009

Hal Daumé III  
hal@cs.utah.edu

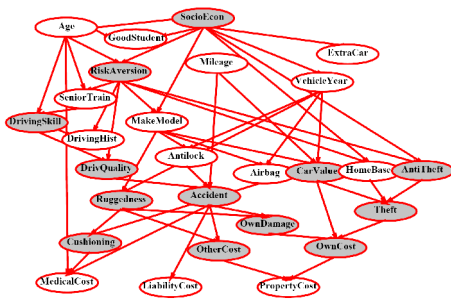
www.cs.utah.edu/~hal/courses/2009S\_AI

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

## Announcements

- Midterm will be handed back Thursday (sorry for delay, but there are a *lot* of them!)
- HW6 is due Thursday
- P4 will be out by Thursday
- Contest timeline:
  - Code out this week
  - Week of April 20th: qualifiers
  - April 28th: Live Berkeley vs Utah Games
- Scott will not be in the lab today

## Example Bayes' Net




## Bayes' Nets

- A Bayes' net is an efficient encoding of a probabilistic model of a domain
- Questions we can ask:
  - Inference: given a fixed BN, what is  $P(X | e)$ ?
  - Representation: given a fixed BN, what kinds of distributions can it encode?
  - Modeling: what BN is most appropriate for a given domain?

## Bayes' Net Semantics

- A Bayes' net:
  - A set of nodes, one per variable  $X$
  - A directed, acyclic graph
  - A conditional distribution of each variable conditioned on its parents (the parameters  $\theta$ )
- Semantics:
  - A BN **defines** a joint probability distribution over its variables:



## Building the (Entire) Joint

- We can take a Bayes' net and build any entry from the full joint distribution it encodes
- Typically, there's no reason to build ALL of it
- We build what we need on the fly
- To emphasize: every BN over a domain **implicitly defines a joint distribution** over that domain, specified by local probabilities and graph structure

UNIVERSITY OF UTAH Hal Daumé III (hal@cs.utah.edu)

## Example: Coin Flips

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

Slide 7 CS 5300: Bayes Nets II

UNIVERSITY OF UTAH Hal Daumé III (hal@cs.utah.edu)

## Example: Traffic

Slide 8 CS 5300: Bayes Nets II

UNIVERSITY OF UTAH Hal Daumé III (hal@cs.utah.edu)

## Example: Alarm Network

B	E	P(A B,E)
T	T	.95
T	F	.94
F	T	.29
F	F	.001

A	P(J A)
T	.90
F	.05

A	P(M A)
T	.70
F	.01

P(B)	
Burglary	.001

P(E)	
Earthquake	.002

Slide 9 CS 5300: Bayes Nets II

UNIVERSITY OF UTAH Hal Daumé III (hal@cs.utah.edu)

## Example: Traffic II

- Variables
- T: Traffic
- R: It rains
- L: Low pressure
- D: Roof drips
- B: Ballgame

Slide 10 CS 5300: Bayes Nets II

UNIVERSITY OF UTAH Hal Daumé III (hal@cs.utah.edu)

## Size of a Bayes' Net

- How big is a joint distribution over N Boolean variables?
- How big is an N-node net if nodes have k parents?
- Both give you the power to calculate
- BNs: Huge space savings!
- Also easier to elicit local CPTs
- Also turns out to be faster to answer queries (next class)

Slide 11 CS 5300: Bayes Nets II

UNIVERSITY OF UTAH Hal Daumé III (hal@cs.utah.edu)

## Bayes' Nets

- So far:
  - What is a Bayes' net?
  - What joint distribution does it encode?
- Next: how to answer queries about that distribution
  - Key idea: conditional independence
  - Last class: assembled BNs using an intuitive notion of conditional independence as causality
  - Today: formalize these ideas
  - Main goal: answer queries about conditional independence and influence
- After that: how to answer numerical queries (inference)

Slide 12 CS 5300: Bayes Nets II

UNIVERSITY OF UTAH

Hal Daumé III (hal@cs.utah.edu)

## Conditional Independence

- Reminder: independence
  - X and Y are **independent** if
- X and Y are **conditionally independent** given Z
- (Conditional) independence is a property of a distribution

Slide 13 CS 5300: Bayes Nets II

UNIVERSITY OF UTAH

Hal Daumé III (hal@cs.utah.edu)

## Example: Independence

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

Slide 14 CS 5300: Bayes Nets II

UNIVERSITY OF UTAH

Hal Daumé III (hal@cs.utah.edu)

## 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

Slide 15 CS 5300: Bayes Nets II

UNIVERSITY OF UTAH

Hal Daumé III (hal@cs.utah.edu)

## Independence in a BN

- Important question about a BN:
  - Are two nodes independent given certain evidence?
  - If yes, can calculate using algebra (really tedious)
  - If no, can prove with a counter example
- Example:
  - Question: are X and Z independent?
  - Answer: *not necessarily*, we've seen examples otherwise: low pressure causes rain which causes traffic.
  - X can influence Z, Z can influence X (via Y)
  - Addendum: they *could* be independent: how?

Slide 16 CS 5300: Bayes Nets II

UNIVERSITY OF UTAH

Hal Daumé III (hal@cs.utah.edu)

## Causal Chains

- This configuration is a "causal chain"
  - X: Low pressure
  - Y: Rain
  - Z: Traffic
- Is X independent of Z given Y?
  - Yes!**
- Evidence along the chain "blocks" the influence

Slide 17 CS 5300: Bayes Nets II

UNIVERSITY OF UTAH

Hal Daumé III (hal@cs.utah.edu)

## Common Cause

- Another basic configuration: two effects of the same cause
  - Y: Project due
  - X: Email busy
  - Z: Lab full
- Are X and Z independent?
  - Yes!**
- Observing the cause blocks influence between effects.

Slide 18 CS 5300: Bayes Nets II

UNIVERSITY OF UTAH

Hal Daumé III (hal@cs.utah.edu)

## Common Effect

- Last configuration: two causes of one effect (v-structures)
- Are X and Z independent?
  - Yes: remember the ballgame and the rain causing traffic, no correlation?
  - Still need to prove they must be (try it!)
- Are X and Z independent given Y?
  - No: remember that seeing traffic put the rain and the ballgame in competition?
- **This is backwards from the other cases**
  - Observing the effect **enables** influence between effects.

X: Raining  
Z: Ballgame  
Y: Traffic

Slide 19 CS 5300: Bayes Nets II

UNIVERSITY OF UTAH

Hal Daumé III (hal@cs.utah.edu)

## 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

Slide 20 CS 5300: Bayes Nets II

UNIVERSITY OF UTAH

Hal Daumé III (hal@cs.utah.edu)

## Reachability

- Recipe: shade evidence nodes
- Attempt 1: if two nodes are connected by an undirected path not blocked by a shaded node, they are conditionally independent
- Almost works, but not quite
  - Where does it break?
  - Answer: the v-structure at T doesn't count as a link in a path unless "active"

Slide 21 CS 5300: Bayes Nets II

UNIVERSITY OF UTAH

Hal Daumé III (hal@cs.utah.edu)

## Reachability (D-Separation)

- Question: Are X and Y conditionally independent given evidence variables {Z}?
- Look for "active paths" from X to Y
- No active paths = independence!
- A path is active if each triple is either a:
  - Causal chain  $A \rightarrow B \rightarrow C$  where B is unobserved (either direction)
  - Common cause  $A \leftarrow B \rightarrow C$  where B is unobserved
  - Common effect (aka v-structure)  $A \rightarrow B \leftarrow C$  where B or one of its descendants is observed

Active Triples Inactive Triples

Slide 23 CS 5300: Bayes Nets II

UNIVERSITY OF UTAH

Hal Daumé III (hal@cs.utah.edu)

## Example

late report

Yes

Slide 24 CS 5300: Bayes Nets II

UNIVERSITY OF UTAH

Hal Daumé III (hal@cs.utah.edu)

## Example

late report

Yes

Yes

Yes

Slide 25 CS 5300: Bayes Nets II

UNIVERSITY OF UTAH

Hal Daumé III (hal@cs.utah.edu)

## Example

- Variables:
  - R: Raining
  - T: Traffic
  - D: Roof drips
  - S: I'm sad
- Questions:

[Redacted] Yes

Slide 26 CS 5300: Bayes Nets II

UNIVERSITY OF UTAH

Hal Daumé III (hal@cs.utah.edu)

## Causality?

- When Bayes' nets reflect the true causal patterns:
  - Often simpler (nodes have fewer parents)
  - Often easier to think about
  - Often easier to elicit from experts
- BNs need not actually be causal
  - Sometimes no causal net exists over the domain
  - E.g. consider the variables *Traffic* and *Drips*
  - End up with arrows that reflect correlation, not causation
- What do the arrows really mean?
  - Topology may happen to encode causal structure
  - Topology only guaranteed to encode conditional independence

Slide 27 CS 5300: Bayes Nets II

UNIVERSITY OF UTAH

Hal Daumé III (hal@cs.utah.edu)

## Example: Traffic

- Basic traffic net
- Let's multiply out the joint

r	1/4
¬r	3/4

r	t	3/4
r	¬t	1/4
¬r	t	1/2
¬r	¬t	1/2

r	t	3/16
r	¬t	1/16
¬r	t	6/16
¬r	¬t	6/16

Slide 28 CS 5300: Bayes Nets II

UNIVERSITY OF UTAH

Hal Daumé III (hal@cs.utah.edu)

## Example: Reverse Traffic

- Reverse causality?

t	9/16
¬t	7/16

r	t	3/16
r	¬t	1/16
¬r	t	6/16
¬r	¬t	6/16

t	r	1/3
t	¬r	2/3
¬t	r	1/7
¬t	¬r	6/7

Slide 29 CS 5300: Bayes Nets II

UNIVERSITY OF UTAH

Hal Daumé III (hal@cs.utah.edu)

## Example: Coins

- Extra arcs don't prevent representing independence, just allow non-independence

h	0.5
t	0.5

h	0.5
t	0.5

h	0.5
t	0.5

h h	0.5
t h	0.5
h t	0.5
t t	0.5

Slide 30 CS 5300: Bayes Nets II

UNIVERSITY OF UTAH

Hal Daumé III (hal@cs.utah.edu)

## Alternate BNs

Slide 31 CS 5300: Bayes Nets II

## Summary

- Bayes nets compactly encode joint distributions
- Guaranteed independencies of distributions can be deduced from BN graph structure
- The Bayes' ball algorithm (aka d-separation)
- A Bayes' net may have other independencies that are not detectable until you inspect its specific distribution