

CS440/ECE448

Lecture 20: Bayesian Networks

Mark Hasegawa-Johnson, 3/2023

Slides are CC0: Public Domain



By Los Angeles - Own work, CC BY-SA 3.0,
<https://commons.wikimedia.org/w/index.php?curid=129209878>

Outline

- Review: Bayesian classifier
- The Los Angeles burglar alarm example
- Bayesian network: A better way to represent knowledge
- Inference using a Bayesian network
- Key ideas: Independence and Conditional independence

Review: Bayesian Classifier

- Class label $Y = y$, drawn from some set of labels
- Observation $X = x$, drawn from some set of features
- Bayesian classifier: choose the class label, y , that minimizes your probability of making a mistake:

$$\hat{y} = \operatorname{argmax}_y P(Y = y | X = x)$$

Today: What if $P(X,Y)$ is complicated, and the naïve Bayes assumption is unreasonable?

- Specifically, suppose that Y is a scalar, but $X = [X_0, \dots, X_{99}]$
- Then, even if every variable is binary, $P(Y = y|X = x)$ is a table with 2^{101} numbers. Hard to learn from data; hard to use.
- The naïve Bayes assumption simplified the problem as

$$P(X_0, \dots, X_{99}, Y) \approx P(Y) \prod_{i=1}^{99} P(X_i|Y)$$

- ... but what if that assumption is unreasonable? Do we then have no alternative besides learning all 2^{101} probabilities?
- Today: an alternative called a Bayesian network

Outline

- Review: Bayesian classifier
- The Los Angeles burglar alarm example
- Bayesian network: A better way to represent knowledge
- Inference using a Bayesian network
- Key ideas: Independence and Conditional independence

The Los Angeles burglar alarm example

- Suppose I have a house in LA. I'm in Champaign.
- My phone beeps in class: I have messages from both of my LA neighbors, John and Mary.
- Does getting messages from both John and Mary mean that my burglar alarm is going off?
- If my burglar alarm is going off, does that mean my house is being robbed, or is it just an earthquake?



By Los Angeles - Own work, CC BY-SA 3.0,
<https://commons.wikimedia.org/w/index.php?curid=129209878>

Variables

- $B = T$ if my house is being burglarized, else $B = F$
- $E = T$ if there's an earthquake in LA right now, else $E = F$
- $A = T$ if my alarm is going off right now, else $A = F$
- $J = T$ if John is texting me, else $J = F$
- $M = T$ if Mary is texting me, else $M = F$

Inference Problem

- Given that $J = T$ and $M = T$, I want to know what is the probability that I'm being burglarized
- In other words, what is $P(B = T | M = T, J = T)$
- How on Earth would I estimate that probability? I don't know how to estimate that.

Available Knowledge

- LA has 1 million houses & 41 burglaries/day: $P(B = T) = \frac{41}{1000000}$
- There are ~20 earthquakes/year: $P(E = T) = \frac{20}{365}$
- My burglar alarm is pretty good:

	$B = F, E = F$	$B = F, E = T$	$B = T, E = F$	$B = T, E = T$
$P(A = T B, E)$	$\frac{1}{100}$	$\frac{99}{100}$	$\frac{99}{100}$	$\frac{99}{100}$

- Mary would text if there was an alarm: $P(M = T|A = T) = \frac{9}{10}$
- ...but she also texts most days anyway: $P(M = T|A = F) = \frac{1}{2}$

Combining the Available Knowledge

Putting it all together, we have ... well, we have a big mess. And that's not including the variable J:

	$B = F$	$B = T$
$P(B, E = F, A = F, M = F)$	$\left(\frac{999959}{1000000}\right) \left(\frac{345}{365}\right) \left(\frac{99}{100}\right) \left(\frac{1}{2}\right)$	$\left(\frac{41}{1000000}\right) \left(\frac{345}{365}\right) \left(\frac{99}{100}\right) \left(\frac{1}{2}\right)$
$P(B, E = F, A = F, M = T)$	$\left(\frac{999959}{1000000}\right) \left(\frac{345}{365}\right) \left(\frac{99}{100}\right) \left(\frac{1}{2}\right)$	$\left(\frac{41}{1000000}\right) \left(\frac{345}{365}\right) \left(\frac{99}{100}\right) \left(\frac{1}{2}\right)$
$P(B, E = F, A = T, M = F)$	$\left(\frac{999959}{1000000}\right) \left(\frac{345}{365}\right) \left(\frac{1}{100}\right) \left(\frac{1}{10}\right)$	$\left(\frac{41}{1000000}\right) \left(\frac{345}{365}\right) \left(\frac{1}{100}\right) \left(\frac{1}{10}\right)$
$P(B, E = F, A = T, M = T)$	$\left(\frac{999959}{1000000}\right) \left(\frac{345}{365}\right) \left(\frac{1}{100}\right) \left(\frac{9}{10}\right)$	$\left(\frac{41}{1000000}\right) \left(\frac{345}{365}\right) \left(\frac{99}{100}\right) \left(\frac{9}{10}\right)$
⋮	⋮	⋮

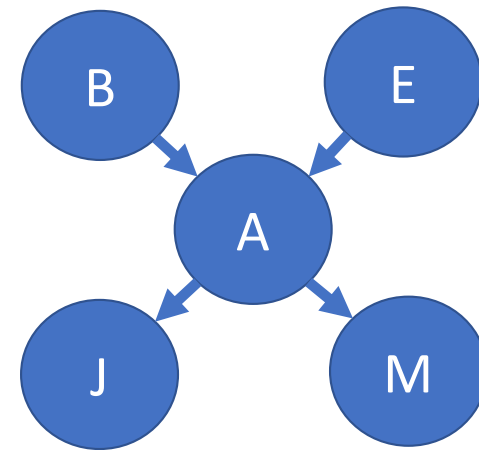
Outline

- Review: Bayesian classifier
- The Los Angeles burglar alarm example
- Bayesian network: A better way to represent knowledge
- Inference using a Bayesian network
- Key ideas: Independence and Conditional independence

Bayesian network: A better way to represent knowledge

A Bayesian network is a graph in which:

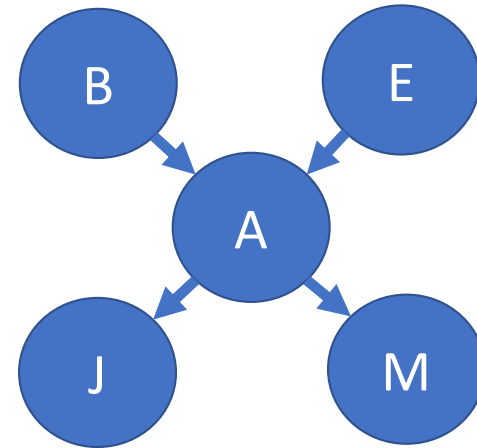
- Each variable is a node.
- An arrow between two nodes means that the child depends on the parent.
- If the child has no direct dependence on the parent, then there is no arrow.



Bayesian network: A better way to represent knowledge

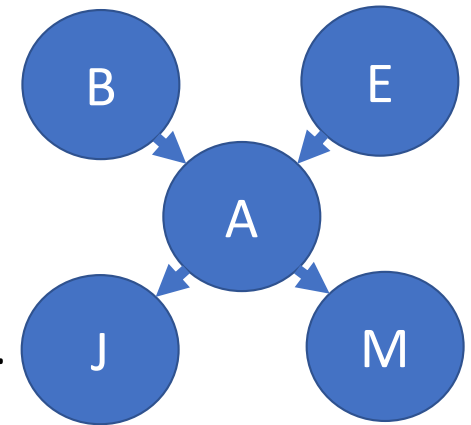
For example, this graph shows my knowledge that:

- My alarm rings if there is a burglary or an earthquake.
- John is more likely to call if my alarm is going off.
- Mary is more likely to call if my alarm is going off.

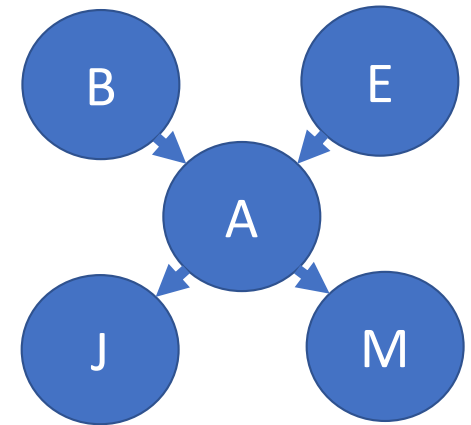


Space complexity

- Without the Bayes network: I have 5 variables, each is binary, so the probability distribution $P(B, E, A, M, J)$ is a table with $2^5 = 32$ entries.
- With the Bayes network:
 - Two of the variables, B and E, depend on nothing else, so I just need to know $P(B = T)$ and $P(E = T)$ --- **1 number** for each of them.
 - M and J depend on A, so I need to know $P(M = T|A = T)$ and $P(M = T|A = F)$ – **2 numbers** for each of them.
 - A depends on both B and E, so I need to know $P(A = T|B = b, E = e)$ for all **4 combinations** of (b, e)
 - Total: $1+1+2+2+4 = 10$ numbers to represent the whole distribution!



Complete description of my knowledge about the burglar alarm



$P(B = T)$	$\frac{41}{1000000}$
------------	----------------------

$P(E = T)$	$\frac{20}{365}$
------------	------------------

	$B = F, E = F$	$B = F, E = T$	$B = T, E = F$	$B = T, E = T$
$P(A = T B, E)$	$\frac{1}{100}$	$\frac{99}{100}$	$\frac{99}{100}$	$\frac{99}{100}$

	$A = F$	$A = T$
$P(J = T A)$	$\frac{1}{8}$	$\frac{3}{4}$

	$A = F$	$A = T$
$P(M = T A)$	$\frac{1}{2}$	$\frac{9}{10}$

Space complexity

- This is a Bayes network to help diagnose problems with your car's audio system.
- Naïve method: 41 binary variables, so the distribution is a table with $2^{41} \approx 2 \times 10^{12}$ entries.
- Bayes network: each variable has at most four parents, so the whole distribution can be described by less than $41 \times 2^4 = 656$ numbers.

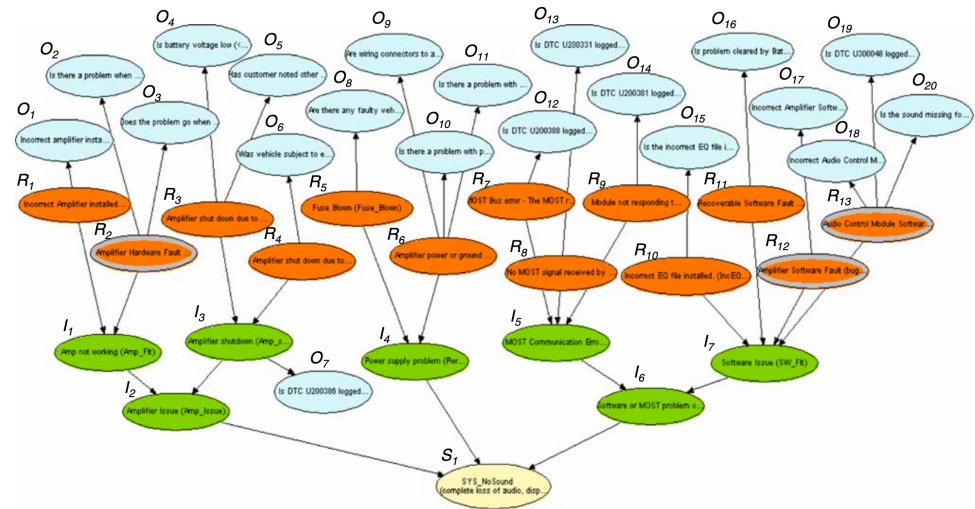


Fig. 6 Bayesian diagnostic model for the symptom “no sound”

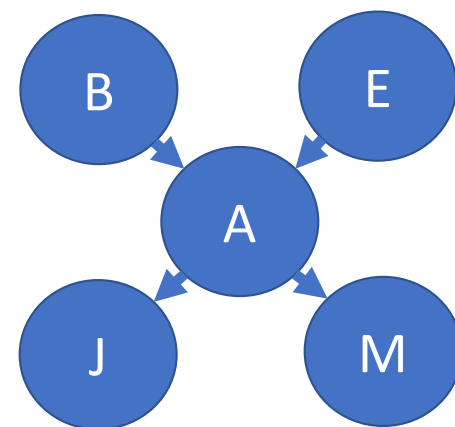
Huang, McMurrin, Dhadyalla & Jones, “Probability-based vehicle fault diagnosis: Bayesian network method,” 2008

Outline

- Review: Bayesian classifier
- The Los Angeles burglar alarm example
- Bayesian network: A better way to represent knowledge
- Inference using a Bayesian network
- Key ideas: Independence and Conditional independence

Inference

Both John and Mary texted me. Am I being burglarized?

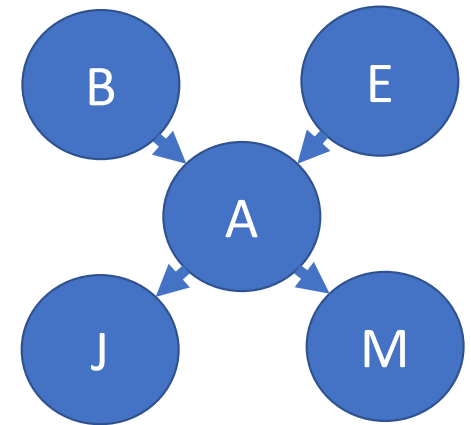


$$P(B = T | J = T, M = T) = \frac{P(B = T, J = T, M = T)}{P(B = T, J = T, M = T) + P(B = F, J = T, M = T)}$$

$$P(B = T, J = T, M = T) = \sum_{e=T}^F \sum_{a=T}^F P(B = T, E = e, A = a, J = T, M = T)$$

$$= \sum_{e=T}^F \sum_{a=T}^F P(B = T)P(E = e)P(A = a | B = T, E = e)P(J = T | A = a)P(M = T | A = a)$$

Time Complexity



- Using a Bayes network doesn't usually change the time complexity of a problem.
- If computing $P(B = T | J = T, M = T)$ required considering 32 possibilities without a Bayes network, it still requires considering 32 possibilities

Some unexpected conclusions

- Burglary is so unlikely that, even if both Mary and John call, it is still more probable that a burglary didn't happen

$$P(B = T | J = T, M = T) < P(B = F | J = T, M = T)$$

- The probability of an earthquake is higher!

$$P(B = T | J = T, M = T) < P(E = T | J = T, M = T)$$

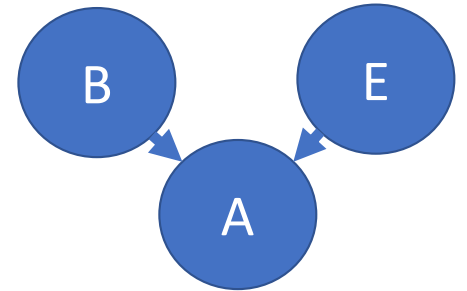
Quiz

- Try the quiz!
- https://us.prairielearn.com/pl/course_instance/129874/assessment/2335333

Outline

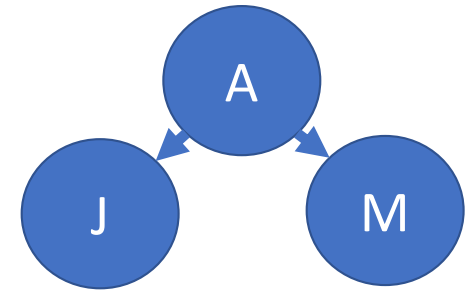
- Review: Bayesian classifier
- The Los Angeles burglar alarm example
- Bayesian network: A better way to represent knowledge
- Inference using a Bayesian network
- **Key ideas: Independence and Conditional independence**

Independence



- The variables B and E are independent
- Days with earthquakes and days w/o earthquakes have the same number of burglaries: $P(B = T|E = T) = P(B = T|E = F) = P(B = T)$.

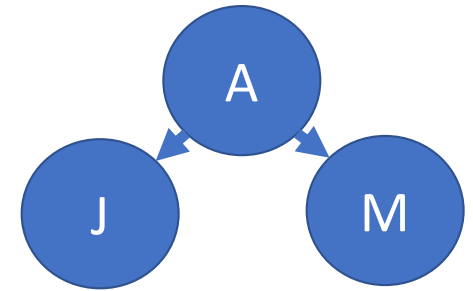
Conditional Independence



- The variables J and M are conditionally independent of one another given knowledge of A
- If you know that there was an alarm, then knowing that John texted gives no extra knowledge about whether Mary will text:

$$P(M = T | J = T, A = T) = P(M = T | J = F, A = T) = P(M = T | A = T)$$

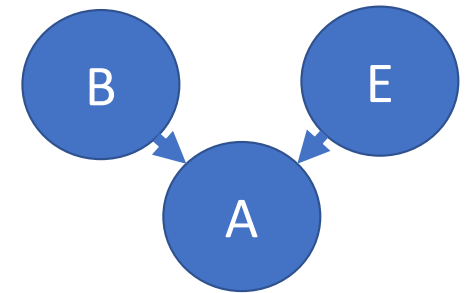
Conditionally Independent variables may not be independent



- The variables J and M are not independent!
- If you know that John texted, that tells you that there was probably an alarm. Knowing that there was an alarm tells you that Mary will probably text you too:

$$P(M = T | J = T) \neq P(M = T | J = F)$$

Independent variables may not be conditionally independent

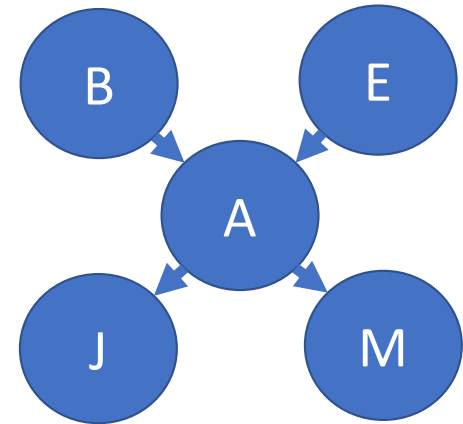


- The variables B and E are not conditionally independent of one another given knowledge of A
- If your alarm is ringing, then you probably have an earthquake OR a burglary. If there is an earthquake, then the conditional probability of a burglary goes down:

$$P(B = T|E = T, A = T) \neq P(B = T|E = F, A = T)$$

- This is called the “explaining away” effect. The earthquake “explains away” the alarm, so you become less worried about a burglary.

How to tell at a glance if variables are independent and/or conditionally independent



- Variables are independent if they have no common ancestors

$$P(B = T, E = T) = P(B = T)P(E = T)$$

- Variables are conditionally independent of one another, given their common ancestors, if (1) they have no common descendants, and (2) none of the descendants of one are ancestors of the other

$$P(J = T, M = T | A = T) = P(J = T | A = T)P(M = T | A = T)$$

Summary

- Review: Bayesian classifier: $\hat{y} = \operatorname{argmax}_y P(Y = y|X = x)$

- Bayesian network: A better way to represent knowledge

- Each variable is a node.
- An arrow between two nodes means that the child depends on the parent.

- Inference using a Bayesian network

$$P(B = T, J = T) = \sum_{e=T}^F \sum_{a=T}^F P(B = T)P(E = e)P(A = a|B = T, E = e)P(J = T|A = a)$$

- Key ideas: Independence and Conditional independence

- Independent = no common ancestors
- Conditionally independent = (1) no common descendants, and (2) none of the descendants of one are ancestors of the other