Warm-up as you walk in

When does a probability table sum to 1?



Announcements

Assignments:

- HW9 (written)
 - Due Tue 4/2, 10 pm

Optional Probability (online)

Midterm:

Mon 4/8, in-class

Course Feedback:

See Piazza post for mid-semester survey

Al: Representation and Problem Solving

Bayes Nets



Instructors: Pat Virtue & Stephanie Rosenthal

Slide credits: CMU AI and http://ai.berkeley.edu

Al-pril Fool's!



Trouble maker credit: Arnav & Pranav

Course Survey

Please fill out on Piazza!



What is the probability of getting a slice with:

- 1) No mushrooms
- 2) Spinach and no mushrooms
- 3) Spinach, when asking for slice with no mushrooms
- Mushrooms
- Spinach
- No spinach
- No spinach and mushrooms
- No spinach when asking for no mushrooms
- No spinach when asking for mushrooms
- Spinach when asking for mushrooms
- No much reasons and no animach



Icons: CC, https://openclipart.org/detail/296791/pizza-slice

You can answer all of these questions:



P(Weather)?

P(Weather | winter)?

P(Weather | winter, hot)?

| Season | Temp | Weather | P(S, T, W) |
|--------|------|---------|------------|
| summer | hot | sun | 0.30 |
| summer | hot | rain | 0.05 |
| summer | cold | sun | 0.10 |
| summer | cold | rain | 0.05 |
| winter | hot | sun | 0.10 |
| winter | hot | rain | 0.05 |
| winter | cold | sun | 0.15 |
| winter | cold | rain | 0.20 |

Two tools to go from joint to query

1. Definition of conditional probability

$$P(A|B) = \frac{P(A,B)}{P(B)}$$

2. Law of total probability (marginalization, summing out)

$$P(A) = \sum_{b} P(A, b)$$

$$P(Y \mid U, V) = \sum_{x} \sum_{z} P(x, Y, z \mid U, V)$$

Two tools to go from joint to query

Joint: $P(H_1, H_2, Q, E)$

Query: $P(Q \mid e)$

1. Definition of conditional probability $P(Q|e) = \frac{P(Q,e)}{P(e)}$

$$P(Q, e) = \sum_{h_1} \sum_{h_2} P(h_1, h_2, Q, e)$$

$$P(e) = \sum_{q} \sum_{h_1} \sum_{h_2} P(h_1, h_2, q, e)$$

P(Weather)?

$$P(Weather | winter)? = \frac{P(W, winter)}{P(W | winter)} = \frac{P(W | winter)}{P(W | win$$

P(Weather | winter, hot)?

| Season | Temp | Weather | P(S, T, W) |
|--------|---------------|---------|------------|
| summer | ummer hot sun | | 0.30 |
| summer | hot | rain | 0.05 |
| summer | cold | sun | 0.10 |
| summer | cold | rain | 0.05 |
| winter | hot | sun | 0.10 |
| winter | hot | rain | 0.05 |
| winter | cold | sun | 0.15 |
| winter | cold | rain | 0.20 |

Joint distributions are the best!

Problems with joints

- Huge
 - *n* variables with *d* values
 - d^n entries
- We aren't given the joint table
 - Usually some set of conditional probability tables



Joint

Build Joint Distribution Using Chain Rule

Conditional Probability Tables and Chain Rule





Build Joint Distribution Using Chain Rule

Two tools to construct joint distribution

1. Product rule

P(A,B) = P(A | B)P(B)P(A,B) = P(B | A)P(A)

2. Chain rule $P(X_1, X_2, ..., X_n) = \prod_i P(X_i \mid X_1, ..., X_{i-1})$ $P(A, B, C) = \underline{P(A)}P(\underline{B} \mid A)P(\underline{C} \mid A, B) \text{ for ordering A, B, C}$ $P(A, B, C) = P(A)P(C \mid A)P(B \mid A, C) \text{ for ordering A, C, B}$

P(A, B, C) = P(C)P(B | C)P(A | C, B) for ordering C, B, A



Process to go from (specific) conditional probability tables to query

- 1. Construct the joint distribution
 - 1. Product Rule or Chain Rule
- 2. Answer query from joint
 - 1. Definition of conditional probability
 - 2. Law of total probability (marginalization, summing out)

Bayes' rule as an example Given: P(E|Q), P(Q) Query: P(Q | e)

- 1. Construct the joint distribution
 - 1. Product Rule or Chain Rule P(E,Q) = P(E|Q)P(Q)
- 2. Answer query from joint
 - 1. Definition of conditional probability

$$P(Q \mid e) = \frac{P(e,Q)}{P(e)}$$

2. Law of total probability (marginalization, summing out) $P(Q \mid e) = \frac{P(e,Q)}{\sum_{a} P(e,q)} \longleftarrow$



Conditional Probability Tables and Chain Rule



Problems

- Huge
 - *n* variables with *d* values
 - *dⁿ* entries
- We aren't given the right tables





P(A) P(B|A) P(C|A) P(D|C) P(E|C)

Query

P⁻

(a)

e)

Bayes Net



P(A) P(B|A) P(C|A) P(D|C) P(E|C)

Build Joint Distribution Using Chain Rule

Chain rule

 $P(X_1, X_2, ..., X_n) = \prod_{i=1}^{n} P(X_i | X_1, ..., X_{i-1})$ $P(C_{1}, C_{2}, C_{3}, C_{4}) =$ $P(C_1)P(C_2|C_1)P(C_3|C_1,C_2)P(C_4|C_1,C_2,C_3)$ $P(h, t, h, t) = \frac{1}{2} \cdot \frac{1}{2} \cdot \frac{1}{2} \cdot \frac{1}{2}$ $P(C_1, C_2, C_3, C_4) = P(C_1) P(C_2) P(C_3) P(C_4)$



Independence



Independence

Two variables X and Y are (absolutely) independent if

• Combine with product rule P(x,y) = P(x|y)P(y) we obtain another form:

 $\forall x, y P(x \mid y) = P(x)$ or $\forall x, y P(y \mid x) = P(y)$

Example: two dice rolls *Roll*₁ and *Roll*₂

 $() \qquad \forall x,y \qquad P(x,y) = P(x) P(y)$

- $P(Roll_1=5, Roll_2=5) = P(Roll_1=5) P(Roll_2=5) = 1/6 \times 1/6 = 1/36$
- $P(Roll_2=5 | Roll_1=5) = P(Roll_2=5)$



Example: Independence

n fair, independent coin flips:





P(Toothache, Cavity, Catch)

If I have a cavity, the probability that the probe catches in it doesn't depend on whether I have a toothache:

P(+catch | +toothache, +cavity) = P(+catch | +cavity)

The same independence holds if I don't have a cavity:

P(+catch | +toothache, -cavity) = P(+catch | -cavity)

Catch is *conditionally independent* of Toothache given Cavity:

P(Catch | Toothache, Cavity) = P(Catch | Cavity)

Equivalent statements:

- P(Toothache | Catch , Cavity) = P(Toothache | Cavity)
- P(Toothache, Catch | Cavity) = P(Toothache | Cavity) P(Catch | Cavity)
- One can be derived from the other easily



Unconditional (absolute) independence very rare (why?)

Conditional independence is our most basic and robust form of knowledge about uncertain environments.

X is conditionally independent of Y given Z if and only if: $\forall x,y,z \quad P(x \mid y, z) = P(x \mid z)$

or, equivalently, if and only if

$$\forall x, y, z \quad P(x, y \mid z) = \underline{P(x \mid z)} P(y \mid z)$$
$$P(x, y \mid z) = \underline{P(x \mid y)} P(y \mid z) P(y \mid z)$$

What about this domain:

- Fire
- Smoke
- Alarm









What about this domain:

- Traffic
- Umbrella
- Raining





Conditional Independence and the Chain Rule

TUR

P(u|R) =

Chain rule:

 $P(x_1, x_2, ..., x_n) = \prod_i P(x_i \mid x_1, ..., x_{i-1})$

P(UIT,R)

Trivial decomposition: P(Rain, Traffic, Umbrella) = P(R) P(T|R)P(U|R,T)

With assumption of conditional independence: P(Rain, Traffic, Umbrella) = P(R) P(T|R) P(u|R)

Conditional Independence and the Chain Rule

Chain rule:

 $P(x_1, x_2, ..., x_n) = \prod_i P(x_i \mid x_1, ..., x_{i-1})$



Trivial decomposition: P(Rain, Traffic, Umbrella) = P(Rain) P(Traffic | Rain) P(Umbrella | Rain, Traffic)

With assumption of conditional independence: *P(Rain, Traffic, Umbrella) = P(Rain) P(Traffic | Rain) P(Umbrella | Rain)*

Bayes nets / graphical models help us express conditional independence assumptions

Bayes'Nets: Big Picture



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





Example Bayes' Net: Insurance



Example Bayes' Net: Car



Graphical Model Notation

Nodes: variables (with domains)

 Can be assigned (observed) or unassigned (unobserved)



- 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!)

P(rav, Catch, ToothAche)

Example: Coin Flips



No interactions between variables: absolute independence

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?

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



Example: Alarm Network

Variables

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





Bayes' Net Semantics



Bayes Nets Syntax Review

One node per random variable DAG

One CPT per node: P(node | *Parents*(node))

Bayes net



Bayes Net Global Semantics



Bayes nets:

Encode joint distributions as product of conditional distributions on each variable

$$P(X_1 \dots X_2) = \prod_i P(X_i | Parents(X_i))$$

Semantics Example



Joint distribution factorization example



• $P(X_1 \dots X_2) = \prod_i P(X_i | Parents(X_i))$

Example: Coin Flips



P(h, h, t, h) =

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

Example: Traffic



P(+r,-t) =





Example: Alarm Network



| E | P(E) | |
|----|-------|--|
| +e | 0.002 | |
| -е | 0.998 | |



| В | Е | Α | P(A B,E) |
|----|----|----|----------|
| +b | +e | +a | 0.95 |
| +b | +e | -a | 0.05 |
| +b | -е | +a | 0.94 |
| +b | -е | -a | 0.06 |
| -b | +e | +a | 0.29 |
| -b | +e | -a | 0.71 |
| -b | -е | +a | 0.001 |
| -b | -е | -a | 0.999 |