CS 112: Modeling Uncertainty in Information Systems

Prof. Jenn Wortman Vaughan

April 4, 2012

Lecture 2

Reminders & Announcements

- Pick up a copy of the course textbook
- Take a look at the course website
- Sign up for Piazza
- Hand in your signed copy of the academic honesty policy

- Homework 1 will be posted by tomorrow…

Today

- Conditional probability
- The multiplication rule & the sequential approach
- Total probability & the divide-and-conquer approach
- Bayes' rule

Probability

- We roll a fair die. What is the probability we roll a 3?

Conditional Probability

- We roll a fair die. What is the probability we roll a 3?
- If we are told that the number rolled is odd, what is the probability that it is 3?

Conditional Probability

- We roll a fair die. What is the probability we roll a 3?
- If we are told that the number rolled is odd, what is the probability that it is 3?
- When all outcomes are equally likely, we can define

$$P(A \mid B) = \#\text{ elements in } A \cap B \,/\, \#\text{ elements in } B$$

Conditional Probability

- We roll a fair die. What is the probability we roll a 3?
- If we are told that the number rolled is odd, what is the probability that it is 3?
- When all outcomes are equally likely, we can define

  $$P(A \mid B) = \# \text{ elements in } A \cap B \, / \, \# \text{ elements in } B$$

- Generalizing this idea, we can define

  $$P(A \mid B) = P(A \cap B) \, / \, P(B)$$

Conditional Probability

Suppose we flip two fair coins.

- What is the probability that the first coin lands on tails given that the second coin lands on tails?

Conditional Probability

Suppose we flip two fair coins.

- What is the probability that the first coin lands on tails given that the second coin lands on tails?
- What is the probability that the first coin lands on tails given that at least one coin lands on tails?

Conditional Probability

Suppose we flip two fair coins.

- What is the probability that the first coin lands on tails given that the second coin lands on tails?
- What is the probability that the first coin lands on tails given that at least one coin lands on tails?

Suppose I purchase one item for a snack. I purchase a cookie with probability 0.5, a brownie with probability 0.25, and an apple with probability 0.25. What is the probability I purchase a cookie given that I purchase a baked good?

A Condition Probability is a Probability Law!

A Condition Probability is a Probability Law!

- We can verify that $P(\cdot|B)$ satisfies the three axioms of probability with universe B

A Condition Probability is a Probability Law!

- We can verify that $P(\cdot|B)$ satisfies the three axioms of probability with universe B
- Any properties that we know are true for probabilities are also true for conditional probabilities!!

The Multiplication Rule

We can express $P(A_1 \cap A_2)$ in terms of a conditional probability:

The Multiplication Rule

We can express $P(A_1 \cap A_2)$ in terms of a conditional probability:

$$P(A_1 \cap A_2) = P(A_1)\, P(A_2 \mid A_1)$$

The Multiplication Rule

We can express $P(A_1 \cap A_2)$ in terms of a conditional probability:

$$P(A_1 \cap A_2) = P(A_1)\, P(A_2 \mid A_1)$$

More generally,

$$P(A_1 \cap A_2 \cap \ldots \cap A_n)$$
$$= P(A_1)\, P(A_2 \mid A_1)\, P(A_3 \mid A_1 \cap A_2) \ldots P(A_n \mid A_1 \cap A_2 \cap \ldots \cap A_{n-1})$$

The Multiplication Rule

If there's an intruder, my alarm will sound with probability 0.99. If there's no intruder, my alarm will sound with probability 0.1. The probability of an intruder is 0.05.

What is the probability of no intruder and my alarm sounding?

The Sequential Approach

1) Represent Ω as a tree that reflects the sequential structure and has the outcomes (i.e., atomic events) as leaves. Internal nodes correspond to non-atomic events, but not all non-atomic events correspond to a single node.

The Sequential Approach

1) Represent Ω as a tree that reflects the sequential structure and has the outcomes (i.e., atomic events) as leaves. Internal nodes correspond to non-atomic events, but not all non-atomic events correspond to a single node.

2) Specify conditional probabilities on branches.

The Sequential Approach

1) Represent Ω as a tree that reflects the sequential structure and has the outcomes (i.e., atomic events) as leaves. Internal nodes correspond to non-atomic events, but not all non-atomic events correspond to a single node.

2) Specify conditional probabilities on branches.

3) Obtain the probability of a leaf by multiplying the probabilities along the corresponding path.

The Sequential Approach

1) Represent Ω as a tree that reflects the sequential structure and has the outcomes (i.e., atomic events) as leaves. Internal nodes correspond to non-atomic events, but not all non-atomic events correspond to a single node.

2) Specify conditional probabilities on branches.

3) Obtain the probability of a leaf by multiplying the probabilities along the corresponding path.

Once you get the hang of it, you don't need to draw the tree…

Uncertainty in LA

Every day I take the freeway to work with probability 1/3; otherwise I take back roads. When I take the freeway, I get stuck in traffic with probability 9/10. When I take the back roads, I get stuck in traffic with probability 1/4.

What is the probability that I take the freeway and get stuck in traffic?

Total Probability Theorem

If $A_1, \ldots, A_n$ partition Ω, then for any event B,

$$P(B) = \sum_{i=1}^{n} P(A_i) P(B \mid A_i)$$

Total Probability Theorem

If $A_1, \ldots, A_n$ partition Ω, then for any event B,

$$P(B) = \sum_{i=1}^{n} P(A_i) P(B \mid A_i)$$

We can prove this using the definition of conditional probability and the additivity axiom…

Total Probability Theorem

If $A_1, \ldots, A_n$ partition Ω, then for any event B,

$$P(B) = \sum_{i=1}^{n} P(A_i) P(B \mid A_i)$$

We can prove this using the definition of conditional probability and the additivity axiom…

This gives us a divide-and-conquer approach to find P(B)

Uncertainty in LA

Every day I take the freeway to work with probability 1/3; otherwise I take back roads. When I take the freeway, I get stuck in traffic with probability 9/10. When I take the back roads, I get stuck in traffic with probability 1/4.

What is the probability that I get stuck in traffic?

Bayes' Rule

For any event A, for any event B with nonzero probability,

$$P(A \mid B) = \frac{P(B \mid A)P(A)}{P(B)}$$

Bayes' Rule

For any event A, for any event B with nonzero probability,

$$P(A \mid B) = \frac{P(B \mid A)P(A)}{P(B)}$$

If A_1, …, A_n partition Ω, then for any *i*,

$$P(A_i \mid B) = \frac{P(B \mid A_i)P(A_i)}{\sum_{j=1}^{n} P(B \mid A_j)P(A_j)}$$

Uncertainty in LA

Every day I take the freeway to work with probability 1/3; otherwise I take back roads. When I take the freeway, I get stuck in traffic with probability 9/10. When I take the back roads, I get stuck in traffic with probability 1/4.

Given that I got stuck in traffic, what is the probability that I took the freeway?

Stress Test

A new stress test can predict whether or not a student is stressed with 80% accuracy; when a student is stressed, it correctly outputs positive with probability 0.8, while if a student is not stressed, it correctly outputs negative with probability 0.8. 95% of students are stressed.

If a student's test results are negative, what is the probability that the student is stressed?