CMPSCI 240: "Reasoning Under Uncertainty"

Lecture 3

Prof. Hanna Wallach
wallach@cs.umass.edu

January 31, 2012

Reminders

- Pick up a copy of B&T
- Check the course website: `http://www.cs.umass.edu/~wallach/courses/s12/cmpsci240/`
- First discussion on Wednesday

Recap

Last Time...

- Probability law: specifies a probability $P(A)$ for all $A \subset \Omega$ either by direct specification or by specification of enough probabilities (either unconditional or conditional) that $P(A)$ can be calculated for any A using these probabilities

Last Time: Discrete Probability Law

- Have unconditional probabilities of atomic events, want to calculate unconditional probability of some non-atomic event

Last Time: Discrete Probability Law

- Have unconditional probabilities of atomic events, want to calculate unconditional probability of some non-atomic event
- e.g., rolling a die: $\Omega = \{1, 2, 3, 4, 5, 6\}$, have $P(i) = 1 \,/\, 6$ for $i = 1 \ldots 6$, want to calculate $P(\text{number is odd})$

Last Time: Discrete Probability Law

- Have unconditional probabilities of atomic events, want to calculate unconditional probability of some non-atomic event
- e.g., rolling a die: $\Omega = \{1, 2, 3, 4, 5, 6\}$, have $P(i) = 1 \,/\, 6$ for $i = 1 \ldots 6$, want to calculate $P(\text{number is odd})$
- Use the discrete probability law: if Ω is discrete and $A = \{x_1, \ldots, x_n\}$ then $P(A) = P(x_1) + \ldots P(x_n)$

Last Time: Probability Axioms etc.

- Have unconditional probabilities of events, want to calculate unconditional probability of some (possibly atomic) event

Last Time: Probability Axioms etc.

- Have unconditional probabilities of events, want to calculate unconditional probability of some (possibly atomic) event
- e.g., $\Omega = \{1, 2, 3, 4, 5, 6\}$, have P(number is odd), P(number $\geq$ 4), and $P(5)$, want to calculate $P(2)$

Last Time: Probability Axioms etc.

- Have unconditional probabilities of events, want to calculate unconditional probability of some (possibly atomic) event
- e.g., $\Omega = \{1, 2, 3, 4, 5, 6\}$, have P(number is odd), P(number $\geq$ 4), and $P(5)$, want to calculate $P(2)$
- Use probability axioms and derivative properties, e.g., $P(A) = 1 - P(A^c)$, $P(A \cup B) \geq P(A) + P(B)$ etc.

Last Time: Conditional Probability

- Have unconditional probabilities of events, want to calculate conditional probability of some (possibly atomic) event

Last Time: Conditional Probability

- Have unconditional probabilities of events, want to calculate conditional probability of some (possibly atomic) event
- e.g., $\Omega = \{1, 2, 3, 4, 5, 6\}$, have P(number is even), P(number ≤ 2), and $P(1)$, want to calculate P(number is $2 \mid$ number is even)

Last Time: Conditional Probability

- Have unconditional probabilities of events, want to calculate conditional probability of some (possibly atomic) event
- e.g., $\Omega = \{1, 2, 3, 4, 5, 6\}$, have P(number is even), P(number ≤ 2), and $P(1)$, want to calculate P(number is 2 | number is even)
- Use conditional probability: $P(A \mid B) = P(A \cap B) \,/\, P(B)$

Last Time: Multiplication Rule

- Have conditional probabilities of events, want to calculate unconditional probability of some (possibly atomic) event

Last Time: Multiplication Rule

- Have conditional probabilities of events, want to calculate unconditional probability of some (possibly atomic) event
- e.g., $\Omega = \{1, 2, 3, 4, 5, 6\}$, have P(number is even) and $P(2 \mid \text{number is even})$, want $P(2)$

Last Time: Multiplication Rule

- Have conditional probabilities of events, want to calculate unconditional probability of some (possibly atomic) event
- e.g., $\Omega = \{1, 2, 3, 4, 5, 6\}$, have $P(\text{number is even})$ and $P(2 \mid \text{number is even})$, want $P(2)$
- Use $P(A \cap B) = P(B)\, P(A \mid B)$ or the multiplication rule: $P(A_1 \cap \ldots \cap A_n) = P(A_1) \ldots P(A_n \mid A_1 \cap \ldots \cap A_{n-1})$

“Sequential” Experiments

“Sequential” Experiments

- Last scenario usually occurs when the experiment has a “sequential” structure, e.g., flipping two coins in a row

"Sequential" Experiments

- Last scenario usually occurs when the experiment has a "sequential" structure, e.g., flipping two coins in a row
- Represent Ω as a tree that reflects the sequential structure and has the outcomes (i.e., atomic events) as leaves

“Sequential” Experiments

- Last scenario usually occurs when the experiment has a “sequential” structure, e.g., flipping two coins in a row
- 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

“Sequential” Experiments

- Last scenario usually occurs when the experiment has a “sequential” structure, e.g., flipping two coins in a row
- 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
- Specify conditional probabilities on branches

Multiplication Rule

- Have sequential experiment, Ω as tree, conditional probabilities on branches, want to calculate unconditional probability of some (possibly atomic) event

Multiplication Rule

- Have sequential experiment, Ω as tree, conditional probabilities on branches, want to calculate unconditional probability of some (possibly atomic) event
- e.g., flipping two coins, $\Omega = \{HH, HT, TH, TT\}$, have $P(\text{first flip is } H)$ and $P(\text{second flip is } H \mid \text{first flip is } H)$, want to calculate $P(HH) = P(\text{both flips are } H)$

Multiplication Rule

- Have sequential experiment, Ω as tree, conditional probabilities on branches, want to calculate unconditional probability of some (possibly atomic) event
- e.g., flipping two coins, $\Omega = \{HH, HT, TH, TT\}$, have $P(\text{first flip is } H)$ and $P(\text{second flip is } H \,|\, \text{first flip is } H)$, want to calculate $P(HH) = P(\text{both flips are } H)$
- Multiply probabilities along path from root to event

Total Probability Theorem and Bayes' Rule

Total Probability Theorem

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

$$\begin{aligned} P(B) &= P(B \cap A_1) + \ldots + P(B \cap A_n) \\ &= \sum_{i=1}^{n} P(A_i)\, P(B \mid A_i) \end{aligned}$$

Total Probability Theorem

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

$$\begin{aligned} P(B) &= P(B \cap A_1) + \ldots + P(B \cap A_n) \\ &= \sum_{i=1}^{n} P(A_i)\, P(B \mid A_i) \end{aligned}$$

- "Divide-and-conquer" approach to finding $P(B)$

Total Probability Theorem

- Have unconditional probabilities of events $A_1, \ldots, A_n$ that partition Ω and conditional probability of B given each A_i for $i = 1, \ldots, n$, want to calculate unconditional probability of B

Total Probability Theorem

- Have unconditional probabilities of events $A_1, \ldots, A_n$ that partition Ω and conditional probability of B given each A_i for $i = 1, \ldots, n$, want to calculate unconditional probability of B
- e.g., $\Omega = \{HH, HT, TH, TT\}$, have $P(\text{first flip is } H)$, $P(\text{second flip is } H \,|\, \text{first flip is } H)$, and $P(\text{second flip is } H \,|\, \text{first flip is } T)$, want $P(\text{second flip is } H)$

Total Probability Theorem

- Have unconditional probabilities of events $A_1, \ldots, A_n$ that partition Ω and conditional probability of B given each A_i for $i = 1, \ldots, n$, want to calculate unconditional probability of B
- e.g., $\Omega = \{HH, HT, TH, TT\}$, have $P(\text{first flip is } H)$, $P(\text{second flip is } H \,|\, \text{first flip is } H)$, and $P(\text{second flip is } H \,|\, \text{first flip is } T)$, want $P(\text{second flip is } H)$
- Use total probabiity theorem: $P(B) = \sum_{i=1}^{n} P(A_i)\, P(B \,|\, A_i)$

Examples of Total Probability Theorem

- e.g., every day I take Route 9 to the department with probability $1/3$; otherwise I take back roads. When I take Route 9, 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?

Rev. Thomas Bayes

Bayes' Rule

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

$$\underbrace{P(A_i \mid B)}_{\text{posterior}} = \frac{P(B \cap A_i)}{P(B)} = \frac{\overbrace{P(A_i)}^{\text{prior}}\, P(B \mid A_i)}{\sum_{i=1}^{n} P(A_i)\, P(B \mid A_i)}$$

Bayes' Rule

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

$$\underbrace{P(A_i \mid B)}_{\text{posterior}} = \frac{P(B \cap A_i)}{P(B)} = \frac{\overbrace{P(A_i)}^{\text{prior}} \, P(B \mid A_i)}{\sum_{i=1}^{n} P(A_i) \, P(B \mid A_i)}$$

- Useful for inference, i.e., where we know $P(B \mid A_i)$ and $P(A_i)$ for every i and want to find $P(A_i \mid B)$ for some i

Examples of Bayes' Rule

- e.g., a test can detect a disease with 95% accuracy: if a person has the disease, the results are positive with probability 0.95, if they don't, the results are negative with probability 0.95. A random person drawn from the population has a probability 0.001 of having the disease. If a person tests positive, what is the probability that they have the disease?

Independence

Independence

- Conditional probability: $P(A \mid B)$ is the probability of event A given that we know the outcome is within event B

Independence

- Conditional probability: $P(A \mid B)$ is the probability of event A given that we know the outcome is within event B
- Suppose event B provides no information about event A?

Independence

- Conditional probability: $P(A \mid B)$ is the probability of event A given that we know the outcome is within event B
- Suppose event B provides no information about event A?
- e.g., flipping two coins: knowing that the first is heads gives no information about whether the second will be heads

Independence

- Conditional probability: $P(A \mid B)$ is the probability of event A given that we know the outcome is within event B

- Suppose event B provides no information about event A?
- e.g., flipping two coins: knowing that the first is heads gives no information about whether the second will be heads

- If the occurrence of event B does not alter the probability of A, then events A and B are independent events

Independence

- Two events A and B are independent if and only if

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

Independence

- Two events A and B are independent if and only if

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

- If $P(B) > 0$ this is equivalent to

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

Independence

- Two events A and B are independent if and only if

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

- If $P(B) > 0$ this is equivalent to

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

- If $P(A) > 0$ this is equivalent to $P(B \,|\, A) = P(B)$

[Got to here in class...]

Examples of Independence

- e.g., are A and A^c independent?

Examples of Independence

- e.g., are A and A^c independent?
- e.g., if A and B are independent, are A and B^c?

Examples of Independence

- e.g., are A and A^c independent?
- e.g., if A and B are independent, are A and B^c?
- e.g., flipping two coins, $H_1 =$ "first flip is heads" $= \{HH, HT\}$, $H_2 =$ is "second flip is heads" $= \{HH, TH\}$, and $D =$ "the flips are different" $= \{HT, TH\}$: H_1 and H_2 are independent, what about H_1 and D?

Independence of Multiple Events

- Events A, B, and C are independent if and only if

$$
\begin{aligned}
P(A \cap B) &= P(A)\,P(B) \\
P(A \cap C) &= P(A)\,P(C) \\
P(B \cap C) &= P(B)\,P(C) \\
P(A \cap B \cap C) &= P(A)\,P(B)\,P(C)
\end{aligned}
$$

Independence of Multiple Events

- Events A, B, and C are independent if and only if

$$
\begin{aligned}
P(A \cap B) &= P(A)\,P(B) \\
P(A \cap C) &= P(A)\,P(C) \\
P(B \cap C) &= P(B)\,P(C) \\
P(A \cap B \cap C) &= P(A)\,P(B)\,P(C)
\end{aligned}
$$

- Pairwise independence does not imply independence

Examples of Independence of Multiple Events

- e.g., flipping two coins, $H_1 =$ "first flip is heads" $= \{HH, HT\}$, $H_2 =$ is "second flip is heads" $= \{HH, TH\}$, and $D =$ "the flips are different" $= \{HT, TH\}$: H_1 and H_2 are independent, H_1 and D are independent, H_2 and D are independent, what about H_1, H_2, and D?

Conditional Independence

- A and B are conditionally independent given C if and only if

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

Conditional Independence

- A and B are conditionally independent given C if and only if

$$P(A \cap B \,|\, C) = P(A \,|\, C)\, P(B \,|\, C)$$

- If $P(B \,|\, C) > 0$ this is equivalent to

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

Conditional Independence

- A and B are conditionally independent given C if and only if

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

- If $P(B \mid C) > 0$ this is equivalent to

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

- If $P(A \mid C) > 0$ this is equivalent to $P(B \mid A \cap C) = P(B \mid C)$

Examples of Conditional Independence

- e.g., flipping two coins, $H_1 =$ "first flip is heads" $= \{HH, HT\}$, $H_2 =$ is "second flip is heads" $= \{HH, TH\}$, and $D =$ "the flips are different" $= \{HT, TH\}$: we already know that H_1 and H_2 are independent, but are they conditionally independent given D?

For Next Time

- Read B&T 1.5, 1.6, 1.7
- Check the course website: `http://www.cs.umass.edu/~wallach/courses/s12/cmpsci240/`