CMPSCI 240: “Reasoning Under Uncertainty”

Lecture 2

Prof. Hanna Wallach
wallach@cs.umass.edu

January 26, 2012

Reminders

- Pick up a copy of B&T
- Check the course website: `http://www.cs.umass.edu/~wallach/courses/s12/cmpsci240/`
- NO cell phones or laptops in class

Recap

Last Time...

- Experiment: a process that results in exactly one of several possible outcomes, e.g., rolling a die

Last Time...

- Experiment: a process that results in exactly one of several possible outcomes, e.g., rolling a die
- Sample space: the set of all possible outcomes of an experiment, e.g., $\Omega = \{1, 2, 3, 4, 5, 6\}$

Last Time...

- Experiment: a process that results in exactly one of several possible outcomes, e.g., rolling a die
- Sample space: the set of all possible outcomes of an experiment, e.g., $\Omega = \{1, 2, 3, 4, 5, 6\}$
- Event: a subset of Ω, e.g., $A =$ "odd number" $= \{1, 3, 5\}$

Last Time...

- Experiment: a process that results in exactly one of several possible outcomes, e.g., rolling a die
- Sample space: the set of all possible outcomes of an experiment, e.g., $\Omega = \{1, 2, 3, 4, 5, 6\}$
- Event: a subset of Ω, e.g., $A =$ "odd number" $= \{1, 3, 5\}$
- Atomic event: event consisting of a single outcome, e.g., $\{1\}$

Last Time...

- Probability law: assigns a probability $P(A)$ to any event $A \subseteq \Omega$ encoding our knowledge or beliefs about the collective "likelihood" of the elements of event A; satisfies 3 axioms

Last Time...

- Probability law: assigns a probability $P(A)$ to any event $A \subseteq \Omega$ encoding our knowledge or beliefs about the collective "likelihood" of the elements of event A; satisfies 3 axioms
- Nonnegativity: $P(A) \geq 0$ for every $A \subseteq \Omega$

Last Time...

- Probability law: assigns a probability $P(A)$ to any event $A \subseteq \Omega$ encoding our knowledge or beliefs about the collective "likelihood" of the elements of event A; satisfies 3 axioms
- Nonnegativity: $P(A) \geq 0$ for every $A \subseteq \Omega$
- Additivity: $P(A \cup B) = P(A) + P(B)$ if A and B are disjoint

Last Time...

- Probability law: assigns a probability $P(A)$ to any event $A \subseteq \Omega$ encoding our knowledge or beliefs about the collective "likelihood" of the elements of event A; satisfies 3 axioms

- Nonnegativity: $P(A) \geq 0$ for every $A \subseteq \Omega$
- Additivity: $P(A \cup B) = P(A) + P(B)$ if A and B are disjoint
- Normalization: $P(\Omega) = 1$

Probabilistic Models

Derivative Properties of Probabilities

- $P(\emptyset) = 0$

Derivative Properties of Probabilities

- $P(\emptyset) = 0$
- $P(A^c) = 1 - P(A)$

Derivative Properties of Probabilities

- $P(\emptyset) = 0$
- $P(A^c) = 1 - P(A)$
- If $A \subset B$, then $P(A) \leq P(B)$

Derivative Properties of Probabilities

- $P(\emptyset) = 0$
- $P(A^c) = 1 - P(A)$
- If $A \subset B$, then $P(A) \leq P(B)$
- $P(A \cup B) = P(A) + P(B) - P(A \cap B)$

Derivative Properties of Probabilities

- $P(\emptyset) = 0$
- $P(A^c) = 1 - P(A)$
- If $A \subset B$, then $P(A) \leq P(B)$
- $P(A \cup B) = P(A) + P(B) - P(A \cap B)$
- $P(A \cup B) \leq P(A) + P(B)$

Derivative Properties of Probabilities

- $P(\emptyset) = 0$
- $P(A^c) = 1 - P(A)$
- If $A \subset B$, then $P(A) \leq P(B)$
- $P(A \cup B) = P(A) + P(B) - P(A \cap B)$
- $P(A \cup B) \leq P(A) + P(B)$
- $P(A \cup B \cup C) = P(A) + P(A^c \cap B) + P(A^c \cap B^c \cap C)$

Discrete Probability Laws

- If Ω is finite, the probability of any event can be derived from the probabilities of the atomic events, i.e., if $A = \{x_1, x_2, \ldots, x_n\}$ then $P(A) = P(x_1) + \ldots + P(x_n)$

Discrete Probability Laws

- If Ω is finite, the probability of any event can be derived from the probabilities of the atomic events, i.e., if $A = \{x_1, x_2, \ldots, x_n\}$ then $P(A) = P(x_1) + \ldots + P(x_n)$
- For a finite sample space, the probabilities of the atomic events completely specify the probability law

Discrete Probability Laws

- If Ω is finite, the probability of any event can be derived from the probabilities of the atomic events, i.e., if $A = \{x_1, x_2, \ldots, x_n\}$ then $P(A) = P(x_1) + \ldots + P(x_n)$
- For a finite sample space, the probabilities of the atomic events completely specify the probability law
- Discrete uniform probability law: if Ω is finite and all outcomes are equally likely, then $P(A) = |A| \,/\, |\Omega|$

Examples of Discrete Probability Laws

- e.g., flipping a coin: what is the probability law?

Examples of Discrete Probability Laws

- e.g., flipping a coin: what is the probability law?
- e.g., rolling two 4-sided dice: $\Omega = \{(1,1),(1,2),\ldots,(4,4)\}$, what is the probability that the sum of the rolls is even? What is the probability that the first roll is equal to the second? What is the probability that at least one roll is a 4?

Conditional Probability

Conditional Probability Laws

- Conditional probability provides a way to reason about the outcome of an experiment using partial information

Conditional Probability Laws

- Conditional probability provides a way to reason about the outcome of an experiment using partial information
- e.g., rolling a die: if we know the number rolled is odd, what is the probability that the number is < 3?

Conditional Probability Laws

- Conditional probability provides a way to reason about the outcome of an experiment using partial information
- e.g., rolling a die: if we know the number rolled is odd, what is the probability that the number is < 3?

- Conditional probability law: assigns a conditional probability $P(A \mid B)$ to any event A encoding our knowledge or beliefs about the collective "likelihood" of the elements of A given that we know the outcome is within event B

Conditional Probability

- Conditional probability of event A given event B:

$$P(A \mid B) = \frac{P(A \cap B)}{P(B)} \text{ assuming } P(B) > 0$$

Conditional Probability

- Conditional probability of event A given event B:

$$P(A \mid B) = \frac{P(A \cap B)}{P(B)} \text{ assuming } P(B) > 0$$

- If all outcomes are equally likely, $P(A \mid B) = |A \cap B| \,/\, |B|$

Conditional Probability

- Conditional probability of event A given event B:

$$P(A \mid B) = \frac{P(A \cap B)}{P(B)} \text{ assuming } P(B) > 0$$

- If all outcomes are equally likely, $P(A \mid B) = |A \cap B| \, / \, |B|$
- All the conditional probability is concentrated on B

Conditional Probability

- Conditional probability of event A given event B:

$$P(A \mid B) = \frac{P(A \cap B)}{P(B)} \text{ assuming } P(B) > 0$$

- If all outcomes are equally likely, $P(A \mid B) = |A \cap B| \,/\, |B|$
- All the conditional probability is concentrated on B
- Conditional probabilities satisfy the probability axioms

Examples of Conditional Probability

- e.g., flipping two coins in a row: what is the probability that the first coin is tails given that at least one coin is tails?

Examples of Conditional Probability

- e.g., flipping two coins in a row: what is the probability that the first coin is tails given that at least one coin is tails?
- e.g., I eat a cookie with probability 0.5, a brownie with probability 0.25, and an apple with probability 0.25: what is the probability I ate a cookie given that I ate a baked good?

Examples of Conditional Probability

- e.g., Competitive eaters Takeru Kobayashi and Joey Chestnut are asked to eat 60 hot dogs in 10 minutes. From past experience, we know that the probability that Takeru is successful is $2/3$, the probability that Joey is successful is $1/3$, and the probability that at least one of them is successful is $3/4$. Assuming that exactly one of them was successful, what is the probability that it was Joey?

Multiplication Rule

- For "sequential" experiments, it can be easier to specify the probability law by specifying conditional probabilities and using them to derive unconditional probabilities:

Multiplication Rule

- For "sequential" experiments, it can be easier to specify the probability law by specifying conditional probabilities and using them to derive unconditional probabilities:

$$
\begin{aligned}
&P(A \cap \ldots \cap B) \\
&\quad = P(A_1)\, P(A_2 \mid A_1) \ldots P(A_n \mid A_1 \cap \ldots \cap A_{n-1})
\end{aligned}
$$

Multiplication Rule

- For "sequential" experiments, it can be easier to specify the probability law by specifying conditional probabilities and using them to derive unconditional probabilities:

$$
\begin{aligned}
&P(A \cap \ldots \cap B) \\
&\quad = P(A_1)\, P(A_2 \mid A_1) \ldots P(A_n \mid A_1 \cap \ldots \cap A_{n-1})
\end{aligned}
$$

- If $n = 2$, this is the definition of conditional probability

Examples of the Multiplication Rule

- e.g., 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?

Multiplication Rule as a Tree

- Form a tree where atomic events are associated with leaves

Multiplication Rule as a Tree

- Form a tree where atomic events are associated with leaves
- Internal nodes are associated with other intersection events

Multiplication Rule as a Tree

- Form a tree where atomic events are associated with leaves
- Internal nodes are associated with other intersection events
- Record conditional probabilities on the branches of the tree

Multiplication Rule as a Tree

- Form a tree where atomic events are associated with leaves
- Internal nodes are associated with other intersection events
- Record conditional probabilities on the branches of the tree

- View the occurrence of an atomic event as the traversal of the branches from the root to the corresponding leaf $\implies$ can obtain the probability of the atomic event by multiplying the conditional probabilities on these branches

For Next Time

- Read B&T 1.4, 1.5
- Check the course website: `http://www.cs.umass.edu/~wallach/courses/s12/cmpsci240/`