CMPSCI 240: "Reasoning Under Uncertainty"

Lecture 6

Prof. Hanna Wallach
wallach@cs.umass.edu

February 9, 2012

Reminders

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

Recap

Last Time: Random Variables

- Random variable: a real-valued function of the outcome of an experiment... though not written as a function!

Last Time: Random Variables

- Random variable: a real-valued function of the outcome of an experiment... though not written as a function!
- A random variable X can be either continuous or discrete, i.e., the range of X can be either continuous or discrete

Last Time: Random Variables

- Random variable: a real-valued function of the outcome of an experiment... though not written as a function!
- A random variable X can be either continuous or discrete, i.e., the range of X can be either continuous or discrete
- Event $\{X = x\}$ is the event consisting of all outcomes in Ω that are mapped to value x by random variable X

Last Time: Probability Mass Functions

- Probability mass function (PMF) of X: denoted by p_X where

$$p_X(x) = P(X = x) = P(\{X = x\})$$

Last Time: Probability Mass Functions

- Probability mass function (PMF) of X: denoted by p_X where

$$p_X(x) = P(X = x) = P(\{X = x\})$$

- If X is a random variable and $g(\cdot)$ is some function, then $Y = g(X)$ is another random variable with PMF

$$p_Y(y) = \sum_{\{x \mid g(x) = y\}} p_X(x)$$

Last Time: Common Discrete Random Variables

- Discrete uniform: X has range $[a, b]$ and PMF

$$p_X(k) = \frac{1}{b-a+1} \text{ for } k = a, a+1, \ldots, b$$

Last Time: Common Discrete Random Variables

- Discrete uniform: X has range $[a, b]$ and PMF

$$p_X(k) = \frac{1}{b-a+1} \text{ for } k = a, a+1, \ldots, b$$

- Bernoulli: X has range $\{0, 1\}$ and PMF

$$p_X(k) = p \text{ if } k = 1 \text{ and } p_X(k) = (1-p) \text{ if } k = 0$$

Last Time: Common Discrete Random Variables

- Binomial: X is used to model the number of successes k in n independent trials and has range $[0, n]$ and PMF

$$p_X(k) = \binom{n}{k} p^k (1-p)^{n-k} \text{ for } k = 0, 1, 2, \ldots, n$$

Last Time: Common Discrete Random Variables

- Binomial: X is used to model the number of successes k in n independent trials and has range $[0, n]$ and PMF

$$p_X(k) = \binom{n}{k} p^k (1-p)^{n-k} \text{ for } k = 0, 1, 2, \ldots, n$$

- e.g., you go to a party with 500 guests. What is the probability that one other guest has the same birthday as you?

Common Discrete Random Variables (Cont.)

Geometric Random Variables

- Suppose we flip a biased coin repeatedly until it lands heads (success). X is a geometric random variable if X is equal to the number of times that we have to flip the coin

Geometric Random Variables

- Suppose we flip a biased coin repeatedly until it lands heads (success). X is a geometric random variable if X is equal to the number of times that we have to flip the coin
- The PMF of a geometric random variable X is

$$p_X(k) = (1-p)^{k-1}\, p \text{ for } k = 1, 2, 3, \ldots$$

Geometric Random Variables

- Suppose we flip a biased coin repeatedly until it lands heads (success). X is a geometric random variable if X is equal to the number of times that we have to flip the coin
- The PMF of a geometric random variable X is

$$p_X(k) = (1-p)^{k-1}\, p \text{ for } k = 1, 2, 3, \ldots$$

- Used to model the number of repeated independent trials up to (and including) the first "successful" trial, e.g., spam

Examples of Geometric Random Variables

- e.g., you have 5 keys for your new apartment but you don't know which key opens the front door. What is the PMF of the number of trials needed to open the front door assuming that at each trial you are equally likely to choose any of the 5 keys?

Poisson Random Variables

- The PMF of a Poisson random variable X is

$$p_X(k) = e^{-\lambda} \frac{\lambda^k}{k!} \text{ for } k = 0, 1, 2, \ldots$$

Poisson Random Variables

- The PMF of a Poisson random variable X is

$$p_X(k) = e^{-\lambda} \frac{\lambda^k}{k!} \text{ for } k = 0, 1, 2, \ldots$$

- A Poisson PMF with $\lambda = np$ is a good approximation for a binomial PMF with very small p and very large n if $k \ll n$

Poisson Random Variables

- The PMF of a Poisson random variable X is

$$p_X(k) = e^{-\lambda} \frac{\lambda^k}{k!} \text{ for } k = 0, 1, 2, \ldots$$

- A Poisson PMF with $\lambda = np$ is a good approximation for a binomial PMF with very small p and very large n if $k \ll n$
- e.g., the number of typos in a book with n words, number of cars (out of n) that crash in a city on a given day, etc.

Examples of Poisson Random Variables

- e.g., you go to a party with 500 guests. What is the probability that one other guest has the same birthday as you?

Expectation

Expectation

- Expected value (or expectation or mean) of X:

$$\mathbb{E}[X] = \sum_{x} x \, p_X(x)$$

Expectation

- Expected value (or expectation or mean) of X:

$$\mathbb{E}[X] = \sum_x x\, p_X(x)$$

- A weighted average of the possible values of X

Expectation

- Expected value (or expectation or mean) of X:

$$\mathbb{E}[X] = \sum_x x\, p_X(x)$$

- A weighted average of the possible values of X
- The value we'd "expect" to get for X "on average" if we repeated the same experiment (and calculated X) many times

Expectation

- Expected value (or expectation or mean) of X:

$$\mathbb{E}[X] = \sum_x x\, p_X(x)$$

- A weighted average of the possible values of X
- The value we'd "expect" to get for X "on average" if we repeated the same experiment (and calculated X) many times
- Useful if we want a single number that "summarizes" p_X

Examples of Expectation

- e.g., on any given day, with probability 0.2, I am not especially tired and will drink only 3 cups of coffee before work. With probability 0.8, however, I am VERY tired and will drink 6 cups of coffee before work. What is the expected number of cups that I will drink on any given day?

Expectations of Standard Random Variables

- e.g., if X is a Bernoulli random variable, what is $\mathbb{E}[X]$?

Expectations of Standard Random Variables

- e.g., if X is a Bernoulli random variable, what is $\mathbb{E}[X]$?
- e.g., if X is a discrete uniform random variable, what is $\mathbb{E}[X]$?

Expectations of Standard Random Variables

- e.g., if X is a Bernoulli random variable, what is $\mathbb{E}[X]$?
- e.g., if X is a discrete uniform random variable, what is $\mathbb{E}[X]$?
- e.g., if X is a Poisson random variable, what is $\mathbb{E}[X]$?

[Got to here in class...]

Expected Value Rule and Linearity of Expectation

- If X is a random variable, $g(\cdot)$ is some linear or nonlinear function, and random variable $Y = g(X)$, what is $\mathbb{E}[Y]$?

Expected Value Rule and Linearity of Expectation

- If X is a random variable, $g(\cdot)$ is some linear or nonlinear function, and random variable $Y = g(X)$, what is $\mathbb{E}[Y]$?

$$\mathbb{E}[Y] = \sum_x g(x)\, p_X(x)$$

Expected Value Rule and Linearity of Expectation

- If X is a random variable, $g(\cdot)$ is some linear or nonlinear function, and random variable $Y = g(X)$, what is $\mathbb{E}[Y]$?

$$\mathbb{E}[Y] = \sum_x g(x)\, p_X(x)$$

- Can compute $\mathbb{E}[Y]$ without computing p_Y directly!

Expected Value Rule and Linearity of Expectation

- If X is a random variable, $g(\cdot)$ is some linear or nonlinear function, and random variable $Y = g(X)$, what is $\mathbb{E}[Y]$?

$$\mathbb{E}[Y] = \sum_x g(x)\, p_X(x)$$

- Can compute $\mathbb{E}[Y]$ without computing p_Y directly!
- e.g., if $Y = aX + b$ for any scalars a and b, what is $\mathbb{E}[Y]$?

Variance

Variance

- Variance of X measures the dispersion of X around $\mathbb{E}[X]$:

$$\mathrm{var}(X) = \underbrace{\mathbb{E}[(X - \mathbb{E}[X])^2]}_{\geq 0} = \mathbb{E}[X^2] - \mathbb{E}[X]^2$$

Variance

- Variance of X measures the dispersion of X around $\mathbb{E}[X]$:

$$\mathrm{var}(X) = \underbrace{\mathbb{E}[(X - \mathbb{E}[X])^2]}_{\geq 0} = \mathbb{E}[X^2] - \mathbb{E}[X]^2$$

- Standard deviation of X is

$$\sigma_X = \sqrt{\mathrm{var}(X)}$$

Examples of Variance

- e.g., on any given day, with probability 0.2, I am not especially tired and will drink only 3 cups of coffee before work. With probability 0.8, however, I am VERY tired and will drink 6 cups of coffee before work. Let X be the number of cups that I will drink. $\mathbb{E}[X] = 5.4$. What is $\text{var}(X)$?

Variance of Standard Random Variables

- e.g., $\mathrm{var}(X)$ and σ_X for a Bernoulli random variable?

Variance of Standard Random Variables

- e.g., $\text{var}(X)$ and σ_X for a Bernoulli random variable?
- e.g., $\text{var}(X)$ and σ_X for a Poisson random variable?

Variance of a Linear Function of a Random Variable

- Let $Y = aX + b$ for any random variable X and scalars a and b. $\mathbb{E}[Y] = a\,\mathbb{E}[X] + b$. What is $\mathrm{var}(Y)$? What about σ_Y?

Multiple Random Variables

Joint PMFs

- If X and Y are discrete random variables associated with the same experiment, then their joint PMF is denoted by $p_{X,Y}$

Joint PMFs

- If X and Y are discrete random variables associated with the same experiment, then their joint PMF is denoted by $p_{X,Y}$
- $p_{X,Y}(x, y)$ is the probability of the event $\{X = x\} \cap \{Y = y\}$:

$$p_{X,Y}(x, y) = P(X = x, Y = y) = P(\{X = x\} \cap \{Y = y\})$$

Joint PMFs

- If X and Y are discrete random variables associated with the same experiment, then their joint PMF is denoted by $p_{X,Y}$
- $p_{X,Y}(x, y)$ is the probability of the event $\{X = x\} \cap \{Y = y\}$:

$$p_{X,Y}(x, y) = P(X = x, Y = y) = P(\{X = x\} \cap \{Y = y\})$$

- e.g., choosing a random faculty member in the department: X is the person's height, Y is the person's weight

Tabular Representation

	y_1	y_2	y_3	y_4
x_1	0.1	0.1	0	0.2
x_2	0.05	0.05	0.1	0
x_3	0	0.1	0.2	0.1

- e.g., $p_{X,Y}(x_2, y_3) = 0.1$, $p_{X,Y}(x_3, y_1) = 0$, ...

Tabular Representation

	y_1	y_2	y_3	y_4
x_1	0.1	0.1	0	0.2
x_2	0.05	0.05	0.1	0
x_3	0	0.1	0.2	0.1

- e.g., $p_{X,Y}(x_2, y_3) = 0.1$, $p_{X,Y}(x_3, y_1) = 0$, ...
- Is this a valid joint PMF for X and Y? How do we know?

Marginal PMFs

	y_1	y_2	y_3	y_4
x_1	0.1	0.1	0	0.2
x_2	0.05	0.05	0.1	0
x_3	0	0.1	0.2	0.1

- We can compute the PMFs of X and Y from $p_{X,Y}$

$$p_X(x) = \sum_y p_{X,Y}(x, y) \text{ and } p_Y(y) = \sum_x p_{X,Y}(x, y)$$

Expected Value Rule and Linearity of Expectation

- If X and Y are discrete random variables and $g(\cdot)$ is some linear or nonlinear function of X and Y, what is $\mathbb{E}[g(X, Y)]$?

Expected Value Rule and Linearity of Expectation

- If X and Y are discrete random variables and $g(\cdot)$ is some linear or nonlinear function of X and Y, what is $\mathbb{E}[g(X, Y)]$?

$$\mathbb{E}[g(X, Y)] = \sum_x \sum_y g(x, y)\, p_{X,Y}(x, y)$$

Expected Value Rule and Linearity of Expectation

- If X and Y are discrete random variables and $g(\cdot)$ is some linear or nonlinear function of X and Y, what is $\mathbb{E}[g(X,Y)]$?

$$\mathbb{E}[g(X,Y)] = \sum_x \sum_y g(x,y)\, p_{X,Y}(x,y)$$

- e.g., what is $\mathbb{E}[aX + bY + c]$ for any scalars a, b, and c?

Three or More Random Variables

- The joint PMF of X, Y, and Z is denoted $p_{X,Y,Z}$ where

$$p_{X,Y,Z}(x,y,z) = P(\{X=x\} \cap \{Y=y\} \cap \{Z=z\})$$

Three or More Random Variables

- The joint PMF of X, Y, and Z is denoted $p_{X,Y,Z}$ where

$$p_{X,Y,Z}(x, y, z) = P(\{X = x\} \cap \{Y = y\} \cap \{Z = z\})$$

- The expected value of $g(X, Y, Z)$ is

$$\mathbb{E}[g(X, Y, Z)] = \sum_x \sum_y \sum_z g(x, y, z)\, P_{X,Y,Z}(x, y, z)$$

Three or More Random Variables

- The joint PMF of X, Y, and Z is denoted $p_{X,Y,Z}$ where

$$p_{X,Y,Z}(x, y, z) = P(\{X = x\} \cap \{Y = y\} \cap \{Z = z\})$$

- The expected value of $g(X, Y, Z)$ is

$$\mathbb{E}[g(X, Y, Z)] = \sum_x \sum_y \sum_z g(x, y, z)\, P_{X,Y,Z}(x, y, z)$$

- $\mathbb{E}[aX + bY + cZ + d] = a\,\mathbb{E}[X] + b\,\mathbb{E}[Y] + c\,\mathbb{E}[Z] + d$

Examples of Three or More Random Variables

- e.g., an office party decides to do a gift exchange for their annual holiday party. Everyone in the office purchases a gift and places it in a box. Each person is then given a random gift from the box. What is the expected number of people who get back the gift that they purchased?

Expectations of Standard Random Variables

- e.g., if X is a binomial random variable, what is $\mathbb{E}[X]$?

For Next Time

- Read B&T 2.6, 2.7, 2.8
- Check the course website: `http://www.cs.umass.edu/~wallach/courses/s12/cmpsci240/`
- First homework is due TOMORROW