


Zoo of Discrete RVs

CSE 312 Summer 21
Lecture 12

Announcements

Real World 1 is due on Wednesday!

Problem Set 4 is due on Thursday!



Facts about Variance



Shifting the variance

We know that

$$\mathbb{E}[aX + c] = a\mathbb{E}[X] + c$$

What happens with variance?

i.e., What is $\text{Var}(X + c)$? What is $\text{Var}(aX)$?

Fill out the poll everywhere so
Kushal knows how long to explain
Go to pollev.com/cse312su21

Facts about Variance

$$\text{Var}(X + c) = \text{Var}(X)$$

Proof:

$$\begin{aligned}\text{Var}(X + c) &= \mathbb{E}[(X + c)^2] - \mathbb{E}[X + c]^2 \\ &= \mathbb{E}[X^2] + \mathbb{E}[2Xc] + \mathbb{E}[c^2] - (\mathbb{E}[X] + c)^2 \\ &= \mathbb{E}[X^2] + 2c\mathbb{E}[X] + c^2 - \mathbb{E}[X]^2 - 2c\mathbb{E}[X] - c^2 \\ &= \mathbb{E}[X^2] - \mathbb{E}[X]^2 \\ &= \text{Var}(X)\end{aligned}$$

Facts about Variance

$$\text{Var}(aX) = a^2 \text{Var}(X)$$

Proof:

$$\begin{aligned} \text{Var}(aX) &= \mathbb{E}[(aX)^2] - (\mathbb{E}[aX])^2 \\ &= a^2 \mathbb{E}[X^2] - (a\mathbb{E}[X])^2 \\ &= a^2 \mathbb{E}[X^2] - a^2 \mathbb{E}[X]^2 \\ &= a^2 (\mathbb{E}[X^2] - \mathbb{E}[X]^2) \\ &= a^2 \text{Var}(X) \end{aligned}$$


Shifting a random variable

For any random variable X , and any constants a, b :

$$\mathbb{E}[aX + b] = a\mathbb{E}[X] + b$$

For any random variable X , and any constants a, b :

$$\text{Var}(aX + b) = a^2\text{Var}(X)$$



Common Distributions

Discrete Random Variable Zoo

There are common **patterns** of experiments:

Flip a [fair/unfair] coin [blah] times and count the number of heads.

Flip a [fair/unfair] coin until the first time that you see a heads

Draw a uniformly random element from [set]

...

Instead of calculating the pmf, cdf, support, expectation, variance,... every time, why not calculate it **once** and look it up every time?


What's our goal?

Your goal is NOT to memorize these facts (it'll be convenient to memorize some of them, but don't waste time making flash cards). Everything is on Wikipedia anyway. I check Wikipedia when I forget these.

Our goals are:

1. Practice expectation, variance, etc. for common distributions we have gotten hints of.
2. Introduce one new distribution we haven't seen at all.
3. Review the first half of the course with some probability calculations.

Zoo!



$X \sim \text{Unif}(a, b)$

$$p_X(k) = \frac{1}{b - a + 1}$$

$$\mathbb{E}[X] = \frac{a + b}{2}$$

$$\text{Var}(X) = \frac{(b - a)(b - a + 2)}{12}$$

$X \sim \text{Ber}(p)$

$$p_X(0) = 1 - p;$$

$$p_X(1) = p$$

$$\mathbb{E}[X] = p$$

$$\text{Var}(X) = p(1 - p)$$

$X \sim \text{Bin}(n, p)$

$$p_X(k) = \binom{n}{k} p^k (1 - p)^{n-k}$$

$$\mathbb{E}[X] = np$$

$$\text{Var}(X) = np(1 - p)$$

$X \sim \text{Geo}(p)$

$$p_X(k) = (1 - p)^{k-1} p$$

$$\mathbb{E}[X] = \frac{1}{p}$$

$$\text{Var}(X) = \frac{1 - p}{p^2}$$

$X \sim \text{NegBin}(r, p)$

$$p_X(k) = \binom{k-1}{r-1} p^r (1-p)^{k-r}$$

$$\mathbb{E}[X] = \frac{r}{p}$$

$$\text{Var}(X) = \frac{r(1-p)}{p^2}$$

$X \sim \text{HypGeo}(N, K, n)$

$$p_X(k) = \frac{\binom{K}{k} \binom{N-K}{n-k}}{\binom{N}{n}}$$

$$\mathbb{E}[X] = n \frac{K}{N}$$


$$\text{Var}(X) = \frac{K(N-K)(N-n)}{N^2(N-1)}$$

$X \sim \text{Poi}(\lambda)$

$$p_X(k) = \frac{\lambda^k e^{-\lambda}}{k!}$$

$$\mathbb{E}[X] = \lambda$$

$$\text{Var}(X) = \lambda$$



Uniform Distribution

Scenario: Uniform

You Roll a Fair Die (or draw a random integer from $1, \dots, n$).

More generally: you want an integer in some range, with each equally likely.

Discrete Uniform Distribution

$$X \sim \text{Unif}(a, b)$$

Parameter a is the minimum value in the support, b is the maximum value in the support.

X is a uniformly random integer between a and b (inclusive)

$$p_X(k) = \frac{1}{b-a+1} \text{ for } k \in \mathbb{Z}, a \leq k \leq b$$

$$F_X(k) = \frac{k-a+1}{b-a+1} \text{ for } k \in \mathbb{Z}, a \leq k \leq b.$$

$$\mathbb{E}[X] = \frac{a+b}{2}$$

$$\text{Var}(X) = \frac{(b-a)(b-a+1)}{12}$$



Bernoulli Distribution

Situation: Bernoulli

You flip a biased coin (once) and want to record whether its heads.

You define an indicator random variable and want to know whether it's 1 or not.

More generally: you have one trial, and some probability p of "success."

Bernoulli Distribution

$$X \sim \text{Ber}(p)$$

Parameter p is probability of success.

X is the indicator random variable that the trial was a success.

$$p_X(0) = 1 - p, p_X(1) = p$$

$$F_X(k) = \begin{cases} 0 & \text{if } k < 0 \\ 1 - p & \text{if } 0 \leq k < 1 \\ 1 & \text{if } k \geq 1 \end{cases}$$

$$\mathbb{E}[X] = p$$


$$\text{Var}(X) = p(1 - p)$$

Some other uses:

Did a particular bit get written correctly on the device?

Did you guess right on a multiple-choice question?

Did a server in a cluster fail?



Binomial Distribution

Situation: Binomial

You flip a coin n times independently, each with a probability p of coming up heads. How many heads are there?

More generally: How many success did you see in n independent trials, where each trial has probability p of success?

Binomial Distribution

$$X \sim \text{Bin}(n, p)$$

n is the number of independent trials.

p is the probability of success for one trial.

X is the number of successes across the n trials.

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

F_X is ugly.

$$\mathbb{E}[X] = np$$

$$\text{Var}(X) = np(1-p)$$

Some other uses:

How many bits were written correctly on the device?

How many questions did you guess right on a multiple-choice test?

How many servers in a cluster failed?

How many keys went to one bucket in a hash table?



Geometric Distribution

Situation: Geometric

You flip a coin (which comes up heads with probability p) until you get a heads. How many flips did you need?

More generally: how many independent trials are needed until the first success?

Geometric Distribution

$$X \sim \text{Geo}(p)$$

p is the probability of success for one trial.

X is the number of trials needed to see the first success.

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

$$F_X(k) = 1 - (1 - p)^k \text{ for } k \in \mathbb{N}$$

$$\mathbb{E}[X] = \frac{1}{p}$$

$$\text{Var}(X) = \frac{1-p}{p^2}$$

Some other uses:

How many bits can we write before one is incorrect?

How many questions do you have to answer until you get one right?

How many times can you run an experiment until it fails for the first time?

Geometric: Expectation

$$\begin{aligned}\mathbb{E}[X] &= \sum_{k=1}^{\infty} k(1-p)^{k-1}p \\ &= p \sum_{k=1}^{\infty} k(1-p)^{k-1} = p \cdot \frac{1}{p^2} = \frac{1}{p}.\end{aligned}$$

$$\text{Var}(X) = \mathbb{E}[X^2] - (\mathbb{E}[X])^2$$

$$\mathbb{E}[X^2] = \sum_{k=1}^{\infty} k^2(1-p)^{k-1}p = p \sum_{k=1}^{\infty} k^2(1-p)^{k-1} = \frac{2-p}{p^2}$$

Geometric Property

Geometric random variables are called “memoryless”

Suppose you’re flipping coins (independently) until you see a heads.
The first three came up tails.

How many flips are *left* until you see the first heads?

It’s another independent copy of the original!

The coin “forgot” it already came up tails 3 times.

Formally...

Let X be the number of flips needed, Y be the flips after the third.

$$\mathbb{P}(Y = k | X \geq 3) = \mathbb{P}(Y = k \cap X \geq 3) / \mathbb{P}(X \geq 3)$$

$$\frac{(1-p)^{k+3-1}p}{(1-p)^3}$$

$$= (1-p)^{k-1}p$$

Which is $p_X(k)$.

Practice

Your music teacher requires you to play a 1000 note song without mistake. You have been practicing, so you have probability of 0.999 of getting each note correct (independent of the others). If you mess up a single note, you must start over and play from the beginning. Let X be the number of times you have to play the song from the start. What is $\mathbb{E}[X]$?

Fill out the poll everywhere so
Kushal knows how long to explain
Go to pollev.com/cse312su21

Practice

Your music teacher requires you to play a 1000 note song without mistake. You have been practicing, so you have probability of 0.999 of getting each note correct (independent of the others). If you mess up a single note, you must start over and play from the beginning. Let X be the number of times you have to play the song from the start. What is $\mathbb{E}[X]$?



Negative Binomial Distribution

Scenario: Negative Binomial

You're playing a carnival game, and there are r little kids nearby who all want a stuffed animal. You can win a single game (and thus win one stuffed animal) with probability p (independently each time) How many times will you need to play the game before every kid gets their toy?

More generally, run independent trials with probability p . How many trials do you need for r successes?



Activity

More generally, run independent trials with probability p . How many trials do you need for r successes?

What's the pmf?

What's the expectation and variance? (hint: linearity)

Fill out the poll everywhere so
Kushal knows how long to explain
Go to pollev.com/cse312su21

Negative Binomial Analysis: PMF

What's the pmf? Well how would we know $X = k$?

Of the first $k - 1$ trials, $r - 1$ must be successes.
And trial k must be a success.

That first part is a lot like a binomial!

It's the $p_Y(r - 1)$ where $Y \sim \text{Bin}(k - 1, p)$

First part gives $\binom{k-1}{r-1}(1-p)^{k-1-(r-1)}p^{r-1} = \binom{k-1}{r-1}(1-p)^{k-r}p^{r-1}$

Second part, multiply by p

Total: $p_X(k) = \binom{k-1}{r-1}(1-p)^{k-r}p^r$

Negative Binomial Analysis: Expectation

What about the expectation?

To see r successes:

We flip until we see success 1.

Then flip until success 2.

... Flip until success r .

The total number of flips is...the sum of geometric random variables!

Negative Binomial Analysis: Expectation

Let Z_1, Z_2, \dots, Z_r be independent copies of $\text{Geo}(p)$

Z_i are called "independent and identically distributed" or "i.i.d."

Because they are independent...and have identical pmfs.

$$X \sim \text{NegBin}(r, p) \quad X = Z_1 + Z_2 + \dots + Z_r.$$

$$\mathbb{E}[X] = \mathbb{E}[Z_1 + Z_2 + \dots + Z_r] = \mathbb{E}[Z_1] + \mathbb{E}[Z_2] + \dots + \mathbb{E}[Z_r] = r \cdot \frac{1}{p}$$

Negative Binomial Analysis: Variance

Let Z_1, Z_2, \dots, Z_r be independent copies of $\text{Geo}(p)$

$$X \sim \text{NegBin}(r, p) \quad X = Z_1 + Z_2 + \dots + Z_r.$$

$$\text{Var}(X) = \text{Var}(Z_1 + Z_2 + \dots + Z_r)$$

Up until now we've just used the observation that $X = Z_1 + \dots + Z_r$.
= $\text{Var}(Z_1) + \text{Var}(Z_2) + \dots + \text{Var}(Z_r)$ because the Z_i are independent.

$$= r \cdot \frac{1-p}{p^2}$$

Negative Binomial

$$X \sim \text{NegBin}(r, p)$$

Parameters: r : the number of successes needed, p the probability of success in a single trial

X is the number of trials needed to get the r^{th} success.

$$p_X(k) = \binom{k-1}{r-1} (1-p)^{k-r} p^r$$

$F_X(k)$ is ugly, don't bother with it.

$$\mathbb{E}[X] = \frac{r}{p}$$

$$\text{Var}(X) = \frac{r(1-p)}{p^2}$$



Hypergeometric Distribution

Scenario: Hypergeometric

You have an urn with N balls, of which K are purple. You are going to draw balls out of the urn **without** replacement.

If you draw out n balls, what is the probability you see k purple ones?

Hypergeometric: Analysis

You have an urn with N balls, of which K are purple. You are going to draw balls out of the urn **without** replacement.

If you draw out n balls, what is the probability you see k purple ones?

Of the K purple, we draw out k choose which k will be drawn

Of the $N - K$ other balls, we will draw out $n - k$, choose which $N - K - (n - k)$ will be removed.

Sample space all subsets of size n

$$\frac{\binom{K}{k} \binom{N-K}{n-k}}{\binom{N}{n}}$$

Hypergeometric: Analysis

$$X = D_1 + D_2 + \cdots + D_n$$

Where D_i is the indicator that draw i is purple.

D_1 is 1 with probability K/N .

What about D_2 ?

$$\mathbb{P}(D_2 = 1) = \frac{K-1}{N-1} \cdot \frac{K}{N} + \frac{K}{N-1} \cdot \frac{K-N}{N} = \frac{K(K-N+K-1)}{N(N-1)} = \frac{K}{N}$$

In general $\mathbb{P}(D_i = 1) = \frac{K}{N}$

It might feel counterintuitive, but it's true!

Hypergeometric: Analysis

$$\mathbb{E}[X]$$

$$= \mathbb{E}[D_1 + \cdots + D_n] = \mathbb{E}[D_1] + \cdots + \mathbb{E}[D_n] = n \cdot \frac{K}{N}$$

Can we do the same for variance?

No! The D_i are dependent. Even if they have the same probability.

Hypergeometric Distribution

$$X \sim \text{HypGeo}(N, K, n)$$

Parameters: A total of N balls in an urn, of which K are successes. Draw n balls without replacement.

X is the number of success balls drawn.

$$p_X(k) = \frac{\binom{K}{k} \binom{N-K}{n-k}}{\binom{N}{n}}$$

$$\mathbb{E}[X] = \frac{nK}{N}$$

$$\text{Var}(X) = n \cdot \frac{K}{N} \cdot \frac{N-K}{N} \cdot \frac{N-n}{N-1}$$