

Week 8

8 July 2021

$$P(a \leq X \leq b) = \int_a^b f_X(x) dx$$

$\underbrace{\hspace{10em}}_{\text{some constant}}$

40 Continuous Uniform Distribution

A random variable X has **continuous uniform distribution on the interval (a, b)** if X has density function $f(x)$ which is constant on the interval (a, b) and 0 everywhere else. The density function is given by:

$$f(x) = \begin{cases} \frac{1}{b-a} & \text{if } a < x < b \\ 0 & \text{otherwise} \end{cases}$$

Why can't we choose any constant? area = 1

In this case, the probability function is super easy! $(a < c < d < b)$

$$P(c \leq X \leq d) = \int_c^d \frac{1}{b-a} dx = \frac{d-c}{b-a}$$

What about the expected value?

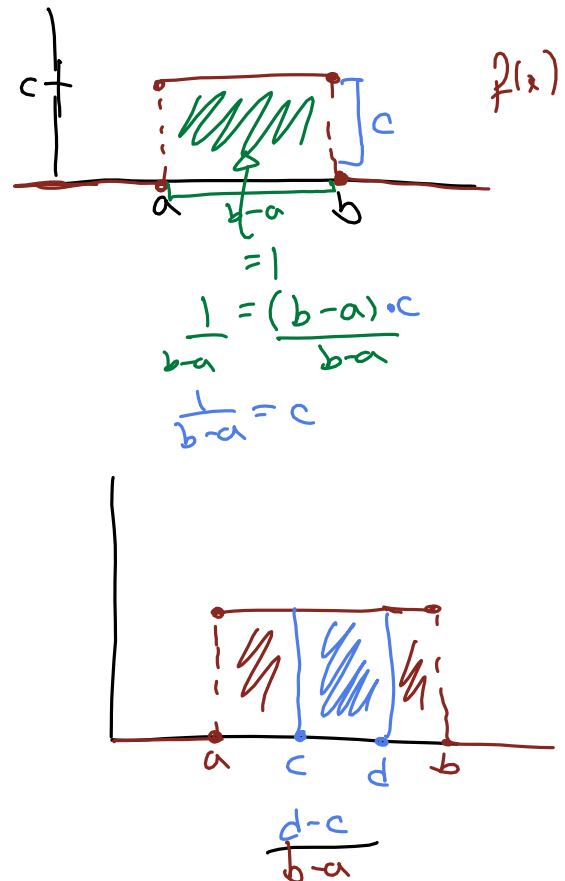
$$\begin{aligned} E(X) &= \int_{-\infty}^{\infty} x f(x) dx \\ &= \int_a^b \frac{x}{b-a} dx \\ &= \left. \frac{x^2}{2(b-a)} \right|_a^b \\ &= \frac{b^2}{2(b-a)} - \frac{a^2}{2(b-a)} = \frac{b^2 - a^2}{2(b-a)} = \boxed{\frac{b+a}{2}} \end{aligned}$$

\swarrow $(b-a)(b+a)$

That's a little complicated, but there is an easier way! We can scale our continuous uniform distribution to a "standard" length. If you recall from the beginning of the semester, we can rescale this to the interval

Wikipedia: [Continuous uniform distribution](#)


Note that the book calls this the "Uniform distribution", but since there is both a discrete and a continuous uniform distribution, we distinguish the two.



$a=0$
 $b=1$

$P(c \leq U \leq d) = \frac{d-c}{b-a} = d-c$

$E(U) = \frac{a+b}{2} = \frac{1}{2}$



std. ct.
Unif.

$f(x) = \begin{cases} 1 & \text{in } 0 < x < 1 \\ 0 & \text{o/w} \end{cases}$

(0, 1). This is called the **Standard continuous uniform distribution**. It is the continuous uniform distribution on (0, 1).

So how do we rescale? Well, first we need to move the entire interval over by a and then we need to scale b down to 1. So we get a standard continuous uniform random variable U where

X is ct.
uniform.

$U = \frac{X-a}{b-a}$

Converting this, we get:

$X = (b-a)U + a$

How does this help?

$$\begin{aligned} E(X) &= E(a + (b-a)U) \\ &= a + (b-a) E(U) \\ &= a + (b-a) \int_0^1 x \cdot 1 dx \\ &= a + (b-a) \left. \frac{x^2}{2} \right|_0^1 \\ &= a + (b-a) \frac{1}{2} \\ &= \frac{2a + b - a}{2} = \frac{a+b}{2} \end{aligned}$$

That integral was much easier to solve! What about the variance? Well, we know that

$E(U^2) = \int_0^1 x^2 \cdot 1 dx = \frac{x^3}{3} \Big|_0^1 = \frac{1}{3}$

So

$$\begin{aligned} \text{Var}(X) &= \text{Var}(a + \underbrace{(b-a)}_{\text{square}} U) \\ &= (b-a)^2 \text{Var}(U) \\ &= (b-a)^2 (E(U^2) - E(U)^2) \\ &= (b-a)^2 \left(\frac{1}{3} - \left(\frac{1}{2}\right)^2 \right) \\ &= (b-a)^2 \left(\frac{4}{12} - \frac{3}{12} \right) \end{aligned}$$

41 Normal Distribution

We've used the normal distribution a lot for approximating the binomial distribution. The *standard normal distribution* is the distribution with density function

$$\phi(z) =$$

Remember that taking the integral of this function is notoriously hard, so instead I'll give you integrals to help out:

$$\int_{-\infty}^{\infty} \phi(z) dz = _ \quad \int_{-\infty}^{\infty} z\phi(z) dz = _ \quad \int_{-\infty}^{\infty} z^2\phi(z) dz = _$$

What this means is that if Z is a standard normal random variable then $E(Z) = 0$ and $\text{Var}(Z) = 1$.

But what if we want our normal formula where we move things left and right and deviate more than normal? (aka the non-standard version) Then we let $X = \mu + \sigma Z$ implying $Z = \frac{X-\mu}{\sigma}$ which should be a formula we've seen before! This gives us a density function

$$\frac{1}{\sigma} \phi\left(\frac{x-\mu}{\sigma}\right) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2}$$

With this density function, we have our *normal distribution*. The expected value is:

$$E(X) =$$

The variance is:

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

The normal distribution is commonly referred to as the *Gaussian distribution*.

Note that since this integral is super difficult, we will continue to use the Φ approximation that we learned in week 3.

Example 41.1 Suppose that we take repeated measurements of the weight of a standard kilogram. Over time, this will slowly decrease in mass! We know that the weight is a normal distribution with expected value 1 kilogram and standard deviation 20 micrograms. What (approximate) proportion of measurements are correct to within 10 micrograms?

Wikipedia: [Normal Distribution](#)

42 Arbitrary continuous distributions

In this section we'll look at two examples of continuous distributions which are not ones that are normally seen. Like this you can see how we can find a distribution from arbitrary density functions and not just the "normal/standard" ones. We do these through examples.

Example 42.1 Suppose that we have a bacterial colony which appears uniformly distributed at random on a circular plate of radius 1. Let R be the random variable which represents the distance from the centre of the plate. What is the probability density, probability (from a to b), expected value and variance of R .


Example 42.2 Let X have the probability density

$$f(x) = \begin{cases} \frac{1}{(1+x)^2} & \text{if } x > 0 \\ 0 & \text{otherwise} \end{cases}$$

What is $P(X > 3)$ and $E(X)$?

43 Empirical distributions

In this section we'll look at what happens when we take empirical data and try and create a continuous distribution out of it. Say I have done some experiment multiple times and I get some random histogram that looks bizarre.



With the above histogram, we might want to draw a “best-fit” curve. This is normally done by something called “numerical analysis” and is a whole topic in mathematics. We won’t get into it to much, but there are ways to approximate this curve. For us, you’ll be given this approximation so don’t worry about it. (Numerical analysis is actually a real cool class!)

The idea is that if we have a random histogram after some experiments, we can ask a numerical analyst to look at the data and create a best-fit function. We can then take that function, say $f(x)$, and create a probability out of it if we add the stipulation that the area under the curve must be 1. This distribution is given by

$$P(a, b) \approx$$

We can actually use indicator functions to help us compute averages from our empirical data. If we let $I_{(a,b)}(x)$ be the indicator function

$$I_{(a,b)}(x) =$$

then we get a way to look at our probabilities. Looking at our histogram we have

$$P(a, b) =$$

where x_i is inside some list of numbers (x_1, x_2, \dots, x_n) from our empirical data. This is basically keeping track of how many of our boxes actually occur. If we combine our two ideas, we end up with

$$P(a, b) \approx$$

This is true since x is 0 outside of (a, b) and so we can make this assertion.

This might seem like a complicated thing to do, because we're basically multiplying by 1, but it actually gives us a powerful tool. What we basically have shown is that

We can generalize this to give us the *integral approximation for averages*. To generalize we just use an arbitrary function $g(x)$ instead of the indicator function.

Notice how both of these functions are giving us the expected value. The left hand side is the expected value of $g(X)$ where X is picked at random from the list (x_1, x_2, \dots, x_n) (aka discrete version). The right hand side is the expected value of $g(X)$ for a random variable X with density $f(x)$.

One thing, right off the bat that we can use this for is moments. If we let $g(x)$ denote the k th moment, *i.e.*, _____, then we have

The left hand side is known as the *k th moment of the empirical distribution*. The right hand side is known as the *k th moment of the theoretical distribution*.

When $k = 1$ and $k = 2$ we have enough information to calculate the expected value and the variance.

But how good is our approximation? To be honest, it's definitely a little complicated. The following is Chebychev's inequality for any $\varepsilon > 0$:

$$P\left(\left|\frac{1}{n}\sum_{i=1}^n g(X_i) - \int_{-\infty}^{\infty} g(x)f(x) dx\right| > \varepsilon\right) \leq \frac{\text{Var}(g(X))}{n\varepsilon^2}$$

I won't solve this or show that it's true. There's more information in the book if you'd like, but the main thing to know here is that we are able to calculate how close the approximation is. Another thing to notice is that if n is very large compared to the variance, then the approximation gets better and better.

Why do we care about all of this? Because integrals are hard! What this is basically telling us is that if we have a difficult integral, we can estimate the integral by a sum! If we take a uniform set of n values, we can approximate our integral using the left hand side. This method is

We never defined the empirical distribution, but you can read more about it on [wiki](#).
Wikipedia: [Empirical Distribution](#)

Wikipedia: [Monte Carlo method](#)

known as the *Monte Carlo method* and is used heavily in physics.

44 Exponential Distribution

Next up, we'll look at distributions where something changes over time. You can think of this as the time it takes for a computer process to be completed, the time it takes for atoms to decay, the time it takes for organisms to evolve, the time it takes for water to boil, etc. Usually time starts from now and goes on into the future. Mathematically we can think of "now" as 0 and the far, far future as ∞ . So, for example, if I want to think about the probability of something happening between 1 and 2 days from now, then if I let T be the random variable which calculates the number of days, then I'm asking for $P(1 \leq T \leq 2)$. Since we're talking about continuous distributions this week, you can expect this to be equal to an integral:

$$P(1 \leq T \leq 2) =$$

If we want to think of something just goes on forever starting from a certain time d , then we can calculate

$$P(d \leq T) =$$

What this means is we can break down $P(a \leq T \leq b)$ into the following

$$P(a \leq T \leq b) =$$

What function we chose to put for $f(t)$ changes based on the model or the question we are asking. There are many different distributions that can be used with respect to time in this way, but the main one we will focus on in this section is the *exponential distribution*.

The *exponential distribution with rate λ* is the continuous distribution where

$$f(t) =$$

Wikipedia: [Exponential distribution](#)

for $t \geq 0$.

What does this distribution look like?

$$P(a \leq T \leq b) =$$

Notice that if $a = 0$ and $b = \infty$, then we get that the entire area is equal to 1.

But look at what this formula implies! If I kept $b = \infty$ and I let a be anything then we have

$$P(a \leq T) =$$

That is so cool! It's actually a super easy formula for us to remember. The nice thing about this is that the expected value and the standard deviation are also really nice numbers:

$$E(T) = \quad \text{and } SD(T) =$$

So the formula is cool and all, but what does this λ mean? λ is basically the value of the instantaneous “bye bye” rate (aka death rate). In other words, it keeps track of how much of a certain thing goes away after a certain amount of time. (Depends on what your “unit” of time is)

Example 44.1 Suppose that you just bought a new charging cable for your phone. You know that, on average, the cable will last you 30 days and you are told that the distribution is exponential. What is the probability that the cable will be working after 15 days?

45 Poisson point process

Remember how the normal/Poisson distribution(s) could be seen as the continuous version(s) of the binomial distribution? It turns out that the exponential distribution is just the continuous version of the geometric distribution! For the geometric distribution we looked over the numbers $\{0, 1, 2, 3, \dots\}$, but in the exponential we look at all non-negative numbers. For a geometric random variable X we had that $E(X) = \frac{1}{p}$ where p is the probability. If our probability is small enough, then $E(X)$ is very large! We can then take this and rescale X by the expected value:

$$\frac{X}{E(X)} =$$

Then, recalling that $P(X > n) = (1-p)^n$ in the geometric case, we have

$$P(pX > t) =$$

This last approximation is coming from the fact that $(1-p) \approx e^{-p}$ for really small p . For example, if $p = 0$:

There's also a way to look at the exponential distribution through the lens of the Poisson distribution. When we had first looked at the binomial distribution, we were keeping track of how many successes we have in n trials. If we had k successes then we had

for the binomial distribution. But, there's another way to think about this. Instead we can look at *when* certain successes occur and look at the gaps in between them.

As a quick example, let's see what I mean about the two ways of looking at this. Suppose we have 10 trials and 2 successes. As an example, we can think of the black dots as successes and the white dots as failures:



- (1) In the first way, we look at it by the number of successes that happen altogether and where these successes might occur. In our example, this is represented by the black circles being in the 1st and the 7th positions. This is where the $\binom{10}{2}$ number is coming from.
- (2) In the second way, we look at the gaps *between* successes. In other words, with our example, we see there are no failures, followed by 1 success, followed by 5 failures, followed by 1 success, followed by 3 failures. So we can represent this as $(0, 5, 3)$. In other words, we can think of this as the ways to add up any three (non-negative) numbers a, b, c such that their sum is $n - k = 10 - 2 = 8$:

$$0 + 5 + 3 = 8$$

Notice that this first way of looking at things is the binomial distribution. On the other hand, looking at things in the second way is the geometric distribution (since we ask for "how long do I have to wait before something happens"). These two ways of looking at things give a description of what's called the *Poisson point process with rate λ* .

Wikipedia: [Poisson point process](#)
The book calls this the "Poisson arrival process".

- (1) **Points/Arrivals:** If we let I be a fixed (time) interval of length t then we can let $N(I)$ be the number of “successes” (usually called *arrivals*). This gives us a Poisson distribution with parameter λt . In this case, λ is the number of successes we expect to see per unit of time.
- (2) **Counts/Times:** Alternatively, we can start from the beginning ($t = 0$) and count the number of success we get as our time increases. This gives us an exponential distribution with expected value $\frac{1}{\lambda}$.

These two ways of looking at things are equivalent.

Example 45.1 The standard example in this case is to look at phone calls, but we’ll use dms instead. Suppose that you get roughly 3 dms per minute (look who’s popular!).

Using the first way of looking at the Poisson point process, we can look at this as a Poisson distribution. Since we’re talking about “per minute” our unit of time is __ minute(s). If we let our interval I be from 2 minutes in the future to 6 minutes in the future, then we know that $t = _$. So then our Poisson distribution parameter is given by _____.

Using the second way of looking at the Poisson point process, we can find an exponential distribution. Since we get 3 dms per minute, that means our average wait time between dms is $\frac{1}{3}$ of a minute. This means, the expected value is $\frac{1}{3} = _$. In other words, we have an exponential distribution with rate _____.

What is the probability that you get 0 dms in the first four minutes ($t = 0$ to $t = 4$)?

What is the probability that we get our first dm after 4 minutes ($t = 4$)?

What is the probability that you get no dms in the first minute and then at most 3 dms in the second minute?

What is the probability that you get no dms in the first minute and the wait time between your second and third dm is more than 2 minutes?

What is the probability that your fifth dm takes more than 4 minutes to arrive?

46 Erlang Distribution

Our next distribution uses the Poisson point process to create a new distribution. If we look at the second definition of the Poisson point process (where we look at the time between successes), then we can keep track of the time of the k th arrival after time 0. Let T_i be the time it takes for a success between the $i - 1$ st success and the i th success, *i.e.*, we assume that T_i has an exponential distribution. So T_1 is how long it takes from 0 until the first success, T_2 is how long it takes from T_1 until second success, etc. Then we let $G_r = T_1 + T_2 + \dots + T_r$, *i.e.*, how long it takes from 0 until the r th success. The distribution $P(G_r)$ is known as the *Erlang distribution with parameters λ and r* . In order to distinguish the two parameters, λ is called the *rate parameter* and r is

Wikipedia: [Erlang distribution](#)

Note that the book calls this a “Gamma distribution”.

known as the *shape parameter*.

The probability density function for an Erlang distribution is given by the function

$$f(x) =$$

Note that it is dependent on the rate λ and the parameter r . Also, note that r *must* be an integer.

Let X be an Erlang random variable with rate λ and parameter r . Then

$$E(X) = \underline{\hspace{2cm}} \quad \text{and} \quad \text{SD}(X) = \underline{\hspace{2cm}}$$

and the tail sum is given by

$$P(X > x) = \sum_{k=0}^{r-1} e^{-\lambda x} \frac{\lambda^k x^k}{k!}$$

Example 46.1 Suppose that it's winter in Toronto and we know that, on average, it snows once every 10 days. Furthermore, suppose that “it snowing” is an exponential distribution. (In other words, the days it snows are exponentially distributed). What is the expected time for the next four snow days to occur?

What is the probability that the next four snow days will occur next week? (7 days from now until 14 days from now)

46.1 Gamma Distribution

In our final distribution for this week we look at the Gamma distribution. If we look at the Erlang distribution we notice that whenever we have G_r we required that r be an integer since it's counting the number of time something happens. If let r be any real number with the same density function as the Erlang distribution, then we get the Gamma distribution. In other words, the *Gamma distribution with parameters λ and r* is the distribution with density function

$$f(x) =$$

Wikipedia: [Gamma distribution](#)

where r is any positive number. In order to distinguish the two parameters, λ is called the *rate parameter* and r is known as the *shape parameter*. Note that if k is an integer, then the Gamma distribution is just the Erlang distribution.