

Section 6: Continuous Random Variables and Central Limit Theorem

Review of Main Concepts

- **Cumulative Distribution Function (cdf):** For any random variable (discrete or continuous) X , the cumulative distribution function is defined as $F_X(x) = \mathbb{P}(X \leq x)$. Notice that this function must be monotonically nondecreasing: if $x < y$ then $F_X(x) \leq F_X(y)$, because $\mathbb{P}(X \leq x) \leq \mathbb{P}(X \leq y)$. Also notice that since probabilities are between 0 and 1, that $0 \leq F_X(x) \leq 1$ for all x , with $\lim_{x \rightarrow -\infty} F_X(x) = 0$ and $\lim_{x \rightarrow +\infty} F_X(x) = 1$.
- **Continuous Random Variable:** A continuous random variable X is one for which its cumulative distribution function $F_X(x) : \mathbb{R} \rightarrow \mathbb{R}$ is continuous everywhere. A continuous random variable has an uncountably infinite number of values.
- **Probability Density Function (pdf or density):** Let X be a continuous random variable. Then the probability density function $f_X(x) : \mathbb{R} \rightarrow \mathbb{R}$ of X is defined as $f_X(x) = \frac{d}{dx}F_X(x)$. Turning this around, it means that $F_X(x) = \mathbb{P}(X \leq x) = \int_{-\infty}^x f_X(t) dt$. From this, it follows that $\mathbb{P}(a \leq X \leq b) = F_X(b) - F_X(a) = \int_a^b f_X(x) dx$ and that $\int_{-\infty}^{\infty} f_X(x) dx = 1$. From the fact that $F_X(x)$ is monotonically nondecreasing it follows that $f_X(x) \geq 0$ for every real number x .

If X is a continuous random variable, note that in general $f_X(a) \neq \mathbb{P}(X = a)$, since $\mathbb{P}(X = a) = F_X(a) - F_X(a) = 0$ for all a . However, the probability that X is close to a is proportional to $f_X(a)$: for small δ , $\mathbb{P}(a - \frac{\delta}{2} < X < a + \frac{\delta}{2}) \approx \delta f_X(a)$.

- **i.i.d. (independent and identically distributed):** Random variables X_1, \dots, X_n are i.i.d. (or iid) if they are independent and have the same probability mass function or probability density function.
- **Discrete to Continuous:**

	Discrete	Continuous
PMF/PDF	$p_X(x) = \mathbb{P}(X = x)$	$f_X(x) \neq \mathbb{P}(X = x) = 0$
CDF	$F_X(x) = \sum_{t < x} p_X(t)$	$F_X(x) = \int_{-\infty}^x f_X(t) dt$
Normalization	$\sum_x p_X(x) = 1$	$\int_{-\infty}^{\infty} f_X(x) dx = 1$
Expectation	$\mathbb{E}[X] = \sum_x x p_X(x)$	$\mathbb{E}[X] = \int_{-\infty}^{\infty} x f_X(x) dx$
LOTUS	$\mathbb{E}[g(X)] = \sum_x g(x) p_X(x)$	$\mathbb{E}[g(X)] = \int_{-\infty}^{\infty} g(x) f_X(x) dx$

- **Standardizing:** Let X be any random variable (discrete or continuous, not necessarily normal), with $\mathbb{E}[X] = \mu$ and $Var(X) = \sigma^2$. If we let $Y = \frac{X - \mu}{\sigma}$, then $\mathbb{E}[Y] = 0$ and $Var(Y) = 1$.
- **Closure of the Normal Distribution:** Let $X \sim \mathcal{N}(\mu, \sigma^2)$. Then, $aX + b \sim \mathcal{N}(a\mu + b, a^2\sigma^2)$. That is, linear transformations of normal random variables are still normal.
- **“Reproductive” Property of Normals:** Let X_1, \dots, X_n be independent normal random variables with $\mathbb{E}[X_i] = \mu_i$ and $Var(X_i) = \sigma_i^2$. Let $a_1, \dots, a_n \in \mathbb{R}$ and $b \in \mathbb{R}$. Then,

$$X = \sum_{i=1}^n (a_i X_i + b) \sim \mathcal{N}\left(\sum_{i=1}^n (a_i \mu_i + b), \sum_{i=1}^n a_i^2 \sigma_i^2\right)$$

There’s nothing special about the parameters – the important result here is that the resulting random variable is still normally distributed.

- **Law of Total Probability (Continuous):** A is an event, and X is a continuous random variable with density function $f_X(x)$.

$$\mathbb{P}(A) = \int_{-\infty}^{\infty} \mathbb{P}(A|X = x) f_X(x) dx$$

- **Central Limit Theorem (CLT):** Let X_1, \dots, X_n be iid random variables with $\mathbb{E}[X_i] = \mu$ and $Var(X_i) = \sigma^2$. Let $X = \sum_{i=1}^n X_i$, which has $\mathbb{E}[X] = n\mu$ and $Var(X) = n\sigma^2$. Let $\bar{X} = \frac{1}{n} \sum_{i=1}^n X_i$, which has $\mathbb{E}[\bar{X}] = \mu$ and $Var(\bar{X}) = \frac{\sigma^2}{n}$. \bar{X} is called the *sample mean*. Then, as $n \rightarrow \infty$, \bar{X} approaches the normal distribution $\mathcal{N}\left(\mu, \frac{\sigma^2}{n}\right)$. Standardizing, this is equivalent to $Y = \frac{\bar{X} - \mu}{\sigma/\sqrt{n}}$ approaching $\mathcal{N}(0, 1)$. Similarly, as $n \rightarrow \infty$, X approaches $\mathcal{N}(n\mu, n\sigma^2)$ and $Y' = \frac{X - n\mu}{\sigma\sqrt{n}}$ approaches $\mathcal{N}(0, 1)$.

It is no surprise that \bar{X} has mean μ and variance σ^2/n – this can be done with simple calculations. The importance of the CLT is that, for large n , regardless of what distribution X_i comes from, \bar{X} is *approximately normally distributed with mean μ and variance σ^2/n* . Don't forget the continuity correction, only when X_1, \dots, X_n are discrete random variables.

- Zoo of Continuous Random Variables

(a) **Uniform:** $X \sim \text{Uniform}(a, b)$ iff X has the following probability density function:

$$f_X(x) = \begin{cases} \frac{1}{b-a} & \text{if } x \in [a, b] \\ 0 & \text{otherwise} \end{cases}$$

$\mathbb{E}[X] = \frac{a+b}{2}$ and $Var(X) = \frac{(b-a)^2}{12}$. This represents each real number from $[a, b]$ to be equally likely.

(b) **Exponential:** $X \sim \text{Exponential}(\lambda)$ iff X has the following probability density function:

$$f_X(x) = \begin{cases} \lambda e^{-\lambda x} & \text{if } x \geq 0 \\ 0 & \text{otherwise} \end{cases}$$

$\mathbb{E}[X] = \frac{1}{\lambda}$ and $Var(X) = \frac{1}{\lambda^2}$. $F_X(x) = 1 - e^{-\lambda x}$ for $x \geq 0$. The exponential random variable is the continuous analog of the geometric random variable: it represents the waiting time to the next event, where $\lambda > 0$ is the average number of events per unit time. Note that the exponential measures how much time passes until the next event (any real number, continuous), whereas the Poisson measures how many events occur in a unit of time (nonnegative integer, discrete). The exponential random variable X is memoryless:

$$\text{for any } s, t \geq 0, \mathbb{P}(X > s + t \mid X > s) = \mathbb{P}(X > t)$$

The geometric random variable also has this property.

(c) **Normal (Gaussian, “bell curve”):** $X \sim \mathcal{N}(\mu, \sigma^2)$ iff X has the following probability density function:

$$f_X(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2} \frac{(x-\mu)^2}{\sigma^2}}, \quad x \in \mathbb{R}$$

$\mathbb{E}[X] = \mu$ and $Var(X) = \sigma^2$. The “standard normal” random variable is typically denoted Z and has mean 0 and variance 1: if $X \sim \mathcal{N}(\mu, \sigma^2)$, then $Z = \frac{X-\mu}{\sigma} \sim \mathcal{N}(0, 1)$. The CDF has no closed form, but we denote the CDF of the standard normal as $\Phi(z) = F_Z(z) = \mathbb{P}(Z \leq z)$. Note from symmetry of the probability density function about $z = 0$ that: $\Phi(-z) = 1 - \Phi(z)$.

- **Closure of the Normal Distribution:** Let $X \sim \mathcal{N}(\mu, \sigma^2)$. Then, $aX + b \sim \mathcal{N}(a\mu + b, a^2\sigma^2)$. That is, linear transformations of normal random variables are still normal.
- **“Reproductive” Property of Normals:** Let X_1, \dots, X_n be independent normal random variables with $\mathbb{E}[X_i] = \mu_i$ and $Var(X_i) = \sigma_i^2$. Let $a_1, \dots, a_n \in \mathbb{R}$ and $b \in \mathbb{R}$. Then,

$$X = \sum_{i=1}^n (a_i X_i + b) \sim \mathcal{N}\left(\sum_{i=1}^n (a_i \mu_i + b), \sum_{i=1}^n a_i^2 \sigma_i^2\right)$$

There's nothing special about the parameters – the important result here is that the resulting random variable is still normally distributed.

- **Law of Total Probability (Continuous):** A is an event, and X is a continuous random variable with density function $f_X(x)$.

$$\mathbb{P}(A) = \int_{-\infty}^{\infty} \mathbb{P}(A \mid X = x) f_X(x) dx$$

- **Central Limit Theorem (CLT):** Let X_1, \dots, X_n be iid random variables with $\mathbb{E}[\sum X_i] = n\mu$ and $\text{Var}(X_i) = \sigma^2$. Let $X = \sum_{i=1}^n X_i$, which has $\mathbb{E}[\sum X_i] = n\mu$ and $\text{Var}(X) = n\sigma^2$. Let $\bar{X} = \frac{1}{n} \sum_{i=1}^n X_i$, which has $\mathbb{E}[\bar{X}] = \mu$ and $\text{Var}(\bar{X}) = \frac{\sigma^2}{n}$. \bar{X} is called the *sample mean*. Then, as $n \rightarrow \infty$, \bar{X} approaches the normal distribution $\mathcal{N}\left(\mu, \frac{\sigma^2}{n}\right)$. Standardizing, this is equivalent to $Y = \frac{\bar{X} - \mu}{\sigma/\sqrt{n}}$ approaching $\mathcal{N}(0, 1)$. Similarly, as $n \rightarrow \infty$, X approaches $\mathcal{N}(n\mu, n\sigma^2)$ and $Y' = \frac{X - n\mu}{\sigma\sqrt{n}}$ approaches $\mathcal{N}(0, 1)$.

It is no surprise that \bar{X} has mean μ and variance σ^2/n – this can be done with simple calculations. The importance of the CLT is that, for large n , regardless of what distribution X_i comes from, \bar{X} is *approximately normally distributed with mean μ and variance σ^2/n* . Don't forget the continuity correction, only when X_1, \dots, X_n are discrete random variables.

1. Uniform2

Robbie decided he wanted to create a “new” type of distribution that will be famous, but he needs some help. He knows he wants it to be continuous and have uniform density, but he needs help working out some of the details. We'll denote a random variable X having the “Uniform-2” distribution as $X \sim \text{Uniform2}(a, b, c, d)$, where $a < b < c < d$. We want the density to be non-zero in $[a, b]$ and $[c, d]$, and zero everywhere else. Anywhere the density is non-zero, it must be equal to the same constant.

- Find the probability density function, $f_X(x)$. Be sure to specify the values it takes on for every point in $(-\infty, \infty)$. (Hint: use a piecewise definition).
- Find the cumulative distribution function, $F_X(x)$. Be sure to specify the values it takes on for every point in $(-\infty, \infty)$. (Hint: use a piecewise definition).

2. Create the distribution

Suppose X is a continuous random variable that is uniform on $[0, 1]$ and uniform on $[1, 2]$, but

$$\mathbb{P}(1 \leq X \leq 2) = 2 \cdot \mathbb{P}(0 \leq X < 1).$$

Outside of $[0, 2]$ the density is 0. What is the PDF and CDF of X ?

3. Max of uniforms

Let U_1, U_2, \dots, U_n be mutually independent Uniform random variables on $(0, 1)$. Find the CDF and pmf for the random variable $Z = \max(U_1, \dots, U_n)$.

4. New PDF?

Alex came up with a function that he thinks could represent a probability density function. He defined the potential pdf for X as $f(x) = \frac{1}{1+x^2}$ defined on $[0, \infty)$. Is this a valid pdf? If not, find a constant c such that the pdf $f_X(x) = \frac{c}{1+x^2}$ is valid. Then find $\mathbb{E}[X]$. (Hints: $\frac{d}{dx}(\tan^{-1} x) = \frac{1}{1+x^2}$, $\tan \frac{\pi}{2} = \infty$, and $\tan 0 = 0$.)

5. Throwing a dart

Consider the closed unit circle of radius r , i.e., $S = \{(x, y) : x^2 + y^2 \leq r^2\}$. Suppose we throw a dart onto this circle and are guaranteed to hit it, but the dart is equally likely to land anywhere in S . Concretely this means that the probability that the dart lands in any particular area of size A (that is entirely inside the circle of radius R), is equal to $\frac{A}{\text{Area of whole circle}}$. The density outside the circle of radius r is 0.

Let X be the distance the dart lands from the center. What is the CDF and pdf of X ? What is $\mathbb{E}[X]$ and $\text{Var}(X)$?

6. A square dartboard?

You throw a dart at an $s \times s$ square dartboard. The goal of this game is to get the dart to land as close to the lower left corner of the dartboard as possible. However, your aim is such that the dart is equally likely to land at any point on the dartboard. Let random variable X be the length of the side of the smallest *square* B in the lower left corner of the dartboard that contains the point where the dart lands. That is, the lower left corner of B must be the same point as the lower left corner of the dartboard, and the dart lands somewhere along the upper or right edge of B . For X , find the CDF, PDF, $\mathbb{E}[X]$, and $Var(X)$.

7. Will the battery last?

Suppose that the number of miles that a car can run before its battery wears out is exponentially distributed with expectation 10,000 miles. If the owner wants to take a 5000 mile road trip, what is the probability that she will be able to complete the trip without replacing the battery, given that the car has already been used for 2000 miles?

8. Batteries and exponential distributions

Let X_1, X_2 be independent exponential random variables, where X_i has parameter λ_i , for $1 \leq i \leq 2$. Let $Y = \min(X_1, X_2)$.

- Show that Y is an exponential random variable with parameter $\lambda = \lambda_1 + \lambda_2$. Hint: Start by computing $\mathbb{P}(Y > y)$. Two random variables with the same CDF have the same pdf. Why?
- What is $Pr(X_1 < X_2)$? (Use the law of total probability.) The law of total probability hasn't been covered in class yet, but will be soon at which point it would be good to revisit this problem!
- You have a digital camera that requires two batteries to operate. You purchase n batteries, labelled $1, 2, \dots, n$, each of which has a lifetime that is exponentially distributed with parameter λ , independently of all other batteries. Initially, you install batteries 1 and 2. Each time a battery fails, you replace it with the lowest-numbered unused battery. At the end of this process, you will be left with just one working battery. What is the expected total time until the end of the process? Justify your answer.
- In the scenario of the previous part, what is the probability that battery i is the last remaining battery as a function of i ? (You might want to use the memoryless property of the exponential distribution that has been discussed.)

9. Grading on a curve

In some classes (not CSE classes) an examination is regarded as being good (in the sense of determining a valid spread for those taking it) if the test scores of those taking it are well approximated by a normal density function. The instructor often uses the test scores to estimate the normal parameters μ and σ^2 and then assigns a letter grade of A to those whose test score is greater than $\mu + \sigma$, B to those whose score is between μ and $\mu + \sigma$, C to those whose score is between $\mu - \sigma$ and μ , D to those whose score is between $\mu - 2\sigma$ and $\mu - \sigma$ and F to those getting a score below $\mu - 2\sigma$. If the instructor does this and a student's grade on the test really is normally distributed with mean μ and variance σ^2 , what is the probability that student will get each of the possible grades A,B,C,D and F?

10. Normal questions

- Let X be a normal random with parameters $\mu = 10$ and $\sigma^2 = 36$. Compute $\mathbb{P}(4 < X < 16)$.
- Let X be a normal random variable with mean 5. If $\mathbb{P}(X > 9) = 0.2$, approximately what is $Var(X)$?

(c) Let X be a normal random variable with mean 12 and variance 4. Find the value of c such that

$$\mathbb{P}(X > c) = 0.10.$$

11. Bad Computer

Each day, the probability your computer crashes is 10%, independent of every other day. Suppose we want to evaluate the computer's performance over the next 100 days.

- (a) Let X be the number of crash-free days in the next 100 days. What distribution does X have? Identify $\mathbb{E}[X]$ and $\text{Var}(X)$ as well. Write an exact (possibly unsimplified) expression for $\mathbb{P}(X \geq 87)$.
- (b) Approximate the probability of at least 87 crash-free days out of the next 100 days using the Central Limit Theorem. Use continuity correction.

Important: continuity correction says that if we are using the normal distribution to approximate

$$\mathbb{P}\left(a \leq \sum_{i=1}^n X_i \leq b\right)$$

where $a \leq b$ are integers and the X_i 's are i.i.d. **discrete** random variables, then, as our approximation, we should use

$$\mathbb{P}(a - 0.5 \leq Y \leq b + 0.5)$$

where Y is the appropriate normal distribution that $\sum_{i=1}^n X_i$ converges to by the Central Limit Theorem.¹

For more details see pages 209-210 in the book.

12. Continuous Law of Total Probability?

This has not been covered in class yet, but will be soon.

In this exercise, we will extend the law of total probability to the continuous case.

- (a) Suppose we flip a coin with probability U of heads, where U is equally likely to be one of $\Omega_U = \{0, \frac{1}{n}, \frac{2}{n}, \dots, 1\}$ (notice this set has size $n + 1$). Let H be the event that the coin comes up heads. What is $\mathbb{P}(H)$?
- (b) Now suppose $U \sim \text{Uniform}(0,1)$ has the *continuous* uniform distribution over the interval $[0, 1]$. Extend the law of total probability to work for this continuous case. (Hint: you may have an integral in your answer instead of a sum).
- (c) Let's generalize the previous result we just used. Suppose E is an event, and X is a continuous random variable with density function $f_X(x)$. Write an expression for $\mathbb{P}(E)$, conditioning on X .

13. Bitcoin users

There is a population of n people. The number of Bitcoin users among these n people is i with probability p_i , where, of course, $\sum_{0 \leq i \leq n} p_i = 1$. We take a random sample of k people from the population (without replacement). Use Bayes Theorem to derive an expression for the probability that there are i Bitcoin users in the population conditioned on the fact that there are j Bitcoin users in the sample. Let B_i be the event that there are i Bitcoin users in the

¹

The intuition here is that, to avoid a mismatch between discrete distributions (whose range is a set of integers) and continuous distributions, we get a better approximation by imagining that a discrete random variable, say W , is a continuous distribution with density function

$$f_W(x) := p_W(i) \quad \text{when } i - 0.5 \leq x < i + 0.5 \text{ and } i \text{ integer}$$

population and let S_j be the event that there are j Bitcoin users in the sample. Your answer should be written in terms of the p_i 's, i, j, n and k .

14. Another continuous r.v.

The density function of X is given by

$$f(x) = \begin{cases} a + bx^2 & \text{when } 0 \leq x \leq 1 \\ 0 & \text{otherwise.} \end{cases}$$

If $E(X) = \frac{3}{5}$, find a and b .

15. Point on a line

A point is chosen at random on a line segment of length L . Interpret this statement and find the probability that the ratio of the shorter to the longer segment is less than $\frac{1}{4}$.

16. Min and max of i.i.d. random variables

Let X_1, X_2, \dots, X_n be i.i.d. random variables each with CDF $F_X(x)$ and pdf $f_X(x)$. Let $Y = \min(X_1, \dots, X_n)$ and let $Z = \max(X_1, \dots, X_n)$. Show how to write the CDF and pdf of Y and Z in terms of the functions $F_X(\cdot)$ and $f_X(\cdot)$.

17. Round off error

Let X be the sum of 100 real numbers, and let Y be the same sum, but with each number rounded to the nearest integer before summing. If the roundoff errors are independent and uniformly distributed between -0.5 and 0.5 , what is the approximate probability that $|X - Y| > 3$?

18. Tweets

A prolific twitter user tweets approximately 350 tweets per week. Let's assume for simplicity that the tweets are independent, and each consists of a uniformly random number of characters between 10 and 140. (Note that this is a discrete uniform distribution.) Thus, the central limit theorem (CLT) implies that the number of characters tweeted by this user is approximately normal with an appropriate mean and variance. Assuming this normal approximation is correct, estimate the probability that this user tweets between 26,000 and 27,000 characters in a particular week. (This is a case where continuity correction will make virtually no difference in the answer, but you should still use it to get into the practice!).

19. Transformations

This has not been covered in class yet and probably won't be. But if you're interested, please read Section 4.4.

Suppose $X \sim \text{Uniform}(0, 1)$ has the continuous uniform distribution on $(0, 1)$. Let $Y = -\frac{1}{\lambda} \log X$ for some $\lambda > 0$.

(a) What is Ω_Y ?

(b) First write down $F_X(x)$ for $x \in (0, 1)$. Then, find $F_Y(y)$ on Ω_Y .

(c) Now find $f_Y(y)$ on Ω_Y (by differentiating $F_Y(y)$ with respect to y . What distribution does Y have?

20. Convolutions

This has not been covered in class. We're not yet sure if we will have time for it, but if you're interested, please read Section 5.5.

Suppose $Z = X + Y$, where $X \perp Y$. (\perp is the symbol for independence. In other words, X and Y are independent.) Z is called the convolution of two random variables. If X, Y, Z are discrete,

$$p_Z(z) = \mathbb{P}(X + Y = z) = \sum_x \mathbb{P}(X = x \cap Y = z - x) = \sum_x p_X(x) p_Y(z - x)$$

If X, Y, Z are continuous,

$$F_Z(z) = \mathbb{P}(X + Y \leq z) = \int_{-\infty}^{\infty} \mathbb{P}(Y \leq z - X | X = x) f_X(x) dx = \int_{-\infty}^{\infty} F_Y(z - x) f_X(x) dx$$

Suppose $X_1 \sim \mathcal{N}(\mu_1, \sigma_1^2)$ and $X_2 \sim \mathcal{N}(\mu_2, \sigma_2^2)$.

- Find an expression for $\mathbb{P}(X_1 < 2X_2)$ using a similar idea to convolution, in terms of $F_{X_1}, F_{X_2}, f_{X_1}, f_{X_2}$. (Your answer will be in the form of a single integral, and requires no calculations – do not evaluate it).
- Find s , where $\Phi(s) = \mathbb{P}(X_1 < 2X_2)$ using the fact that linear combinations of independent normal random variables are still normal.