W12 Notes

Summations

01 Theory

In many contexts it is useful to consider random variables that are summations of a large number of variables.

Summation formulas: $E[X]$ and $\mathrm{Var}[X]$

Suppose X is a large sum of random variables:

$$X = X_1 + X_2 + \cdots + X_n$$

Then:

$$E[X] = E[X_1] + E[X_2] + \cdots + E[X_n]$$

$$\mathrm{Var}[X] = \mathrm{Var}[X_1] + \cdots + \mathrm{Var}[X_n] + 2\sum_{i<j} \mathrm{Cov}[X_i, X_j]$$

If X_i and X_j are uncorrelated (e.g. if they are independent):

$$\mathrm{Var}[X] = \mathrm{Var}[X_1] + \cdots + \mathrm{Var}[X_n]$$

Extra - Derivation of variance of a sum

Using the definition:

$$\begin{aligned}
\mathrm{Var}[X_1 + \cdots + X_n] &= E\Big[(X_1 + \cdots + X_n - (\mu_{X_1} + \cdots + \mu_{X_n}))^2\Big] \\
&= E\Big[\big((X_1 - \mu_{X_1}) + \cdots + (X_n - \mu_{X_n})\big)^2\Big] \\
&= E\left[\sum_{i,j}(X_i - \mu_{X_i})(X_j - \mu_{X_j})\right] \\
&= \sum_{i,j} \mathrm{Cov}(X_i, X_j) \\
&= \sum_i \mathrm{Var}[X_i] + 2\sum_{i<j} \mathrm{Cov}[X_i, X_j]
\end{aligned}$$

In the last line we use the fact that $\mathrm{Cov}[X, X] = \mathrm{Var}[X]$ for the first term, and the symmetry property of covariance for the second term with the factor of 2.

02 Illustration

Example - Binomial expectation and variance

Suppose we have repeated Bernoulli trials $X_1, \ldots, X_n$ with $X_i \sim \mathrm{Ber}(p)$.

The sum is a binomial variable: $S_n = \sum_{i=1}^{n} X_i$.

We know $E[X_i] = p$ and $\text{Var}[X_i] = pq$.

The summation rule for expectation:

$$E[S_n] \;=\; \sum_{i=1}^{n} E[X_i] \quad \gg\gg \quad \sum_{i=1}^{n} p \quad \gg\gg \quad np$$

The summation rule for variance:

$$\text{Var}[S_n] \;=\; \sum_{i=1}^{n} \text{Var}[X_i] + 2\sum_{i<j} \text{Cov}[X_i, X_j]$$

$$\gg\gg \quad \sum_{i=1}^{n} pq + 2 \cdot 0 \quad \gg\gg \quad npq$$

Example - Pascal expectation and variance

(1) Let $X \sim \text{Pasc}(\ell, p)$.

Let X_1, X_2, ... be independent random variables, where:

- X_1 counts the trials until the first success
- X_2 counts the trials *after* the first success until the *second* success
- X_i counts the trials after the $(i-1)^{th}$ success until the i^{th} success

Observe that $X = \sum_{i=1}^{\ell} X_i$.

---

(2) Notice that $X_i \sim \text{Geom}(p)$ for every i. Therefore:

$$E[X_i] \;=\; \frac{1}{p} \qquad \text{Var}[X_i] \;=\; \frac{1-p}{p^2}$$

---

(3) Using the summation rule, conclude:

$$E[X] \quad \gg\gg \quad \sum_{i=1}^{\ell} \frac{1}{p} \quad \gg\gg \quad \frac{\ell}{p}$$

$$\text{Var}[X] \quad \gg\gg \quad \sum_{i=1}^{\ell} \frac{q}{p^2} \quad \gg\gg \quad \frac{\ell q}{p^2}$$

Example - Multinomial covariances

Each trial of an experiment has possible outcomes labeled $1, \ldots, r$ with probabilities of occurrence $p_1, \ldots, p_r$. The experiment is run n times.

Let X_i count the number of occurrences of outcome i. So $X_i \sim \text{Bin}(n, p_i)$.

Find $\text{Cov}[X_i, X_j]$.

Solution

Notice that $X_i + X_j$ is also a binomial variable with success probability $p = p_i + p_j$.
('Success' is an outcome of either i or j. 'Failure' is any other value.)

The variance of a binomial is known to be npq.

Compute $\text{Cov}[X_i, X_j]$ by solving:

$$\begin{aligned} \text{Var}[X_i + X_j] &= \text{Var}[X_i] + \text{Var}[X_j] + 2\text{Cov}[X_i, X_j] \\ n(p_i + p_j)(1 - (p_i + p_j)) &= np_i(1 - p_i) + np_j(1 - p_j) + 2\text{Cov}[X_i, X_j] \\ \gg\gg \quad \text{Cov}[X_i, X_j] &= \boxed{-np_ip_j} \end{aligned}$$

Months with a birthday

Suppose study groups of 10 are formed from a large population.

For a typical study group, how many months out of the year contain a birthday of a member of the group? (Assume all 12 months have equal duration.)

Solution

(1) Let X_i be 1 if month i contains a birthday, and 0 otherwise.

So we seek $E[X_1 + \cdots + X_{12}]$. This equals $E[X_1] + \cdots + E[X_{12}]$.

The answer will be $12E[X_i]$ because all terms are equal.

---

(2) For a given i:

$$P[\text{no birthday in month } i] \quad = \quad \left(\frac{11}{12}\right)^{10}$$

The complement event:

$$P[\text{at least one birthday in month } i] \quad = \quad 1 - \left(\frac{11}{12}\right)^{10}$$

---

(3) Therefore:

$$12E[X_i] \;=\; 12\left(1 - \left(\frac{11}{12}\right)^{10}\right) \quad \gg\gg \quad \boxed{6.97}$$

Example - Hats in the air

All n sailors throws their hats in the air, and catch a random hat when they fall back down.

(a) How many sailors do you expect will catch the hat they own?

(b) What is the variance of this number?

Solution

Strangely, the answers are both 1, regardless of the number of sailors. Here is the reasoning:

(a)
Let $X_i = 1$ when sailor i catches their own hat, and $X_i = 0$ otherwise. Thus X_i is Bernoulli with $p = 1/n$.

Now $X = \sum_{i=1}^{n} X_i$ counts the total number of hats caught by their owners.

Note that $E[X_i] = 1/n$. Therefore:

$$E[X] \quad \gg\gg \quad E\left[\sum_{i=1}^{n} X_i\right]$$

$$\gg\gg \quad \sum_{i=1}^{n} E[X_i] \quad \gg\gg \quad \sum_{i=1}^{n} \frac{1}{n} \quad \gg\gg \quad 1$$

(b)
We know:

$$\text{Var}[X] \;=\; \sum_{i=1}^{n} \text{Var}[X_i] + 2\sum_{i<j} \text{Cov}[X_i, X_j]$$

Now calculate $\text{Var}[X_i]$:

Use $\text{Var}[X_i] \;=\; E[X_i^2] - E[X_i]^2$. Observe that $X_i^2 = X_i$. Therefore:

$$\text{Var}[X_i] \quad \gg\gg \quad \frac{1}{n} - \frac{1}{n^2} \quad \gg\gg \quad \frac{n-1}{n^2}$$

Now calculate $\text{Cov}[X_i, X_j]$:

$$\text{Cov}[X_i, X_j] \;=\; E[X_i X_j] - E[X_i]E[X_j]$$

We need to compute $E[X_i X_j]$.

Notice that $X_i X_j = 1$ when i and j *both* catch their own hats, and 0 otherwise. So it is Bernoulli. Then:

$$P[X_i = 1 \text{ and } X_j = 1] \quad \gg\gg \quad \frac{1}{n(n-1)}$$

$$\gg\gg \quad E[X_i X_j] \;=\; \frac{1}{n(n-1)}$$

Therefore:

$$\text{Cov}[X_i, X_j] \quad \gg\gg \quad \frac{1}{n(n-1)} - \frac{1}{n} \cdot \frac{1}{n}$$

$$\gg\gg \quad \frac{1}{n^2(n-1)}$$

Putting everything together:

$$\text{Var}[X] \quad \gg\gg \quad \sum_{i=1}^{n} \text{Var}[X_i] + 2 \sum_{i<j} \text{Cov}[X_i, X_j]$$

$$\gg\gg \quad \sum_{i=1}^{n} \frac{n-1}{n^2} + 2 \sum_{i<j} \frac{1}{n^2(n-1)}$$

$$\gg\gg \quad \frac{n-1}{n} + n(n-1)\frac{1}{n^2(n-1)} \quad \gg\gg \quad 1$$

Central Limit Theorem

03 Theory

Video by 3Blue1Brown:

- Central limit theorem

IID variables

Random variables are called **independent, identically distributed** when they are independent and have the same distribution.

⚠ IID variables: Same distribution, different values

Independent variables cannot be correlated, so the values taken by IID variables will disagree on all (most) outcomes.

We do have:

$$\text{same distribution} \quad \iff \quad \text{same PMF or PDF}$$

Standardization

Suppose X is any random variable.

The **standardization** of X is:

$$Z \;=\; \frac{X - \mu_X}{\sigma_X}$$

The variable Z has $E[Z] = 0$ and $\text{Var}[Z] = 1$. We can reconstruct X by:

$$X \;=\; \sigma_X Z + \mu_X$$

Suppose X_1, X_2, …, X_n is a collection of IID random variables.

Define:

$$S_n = \sum_{i=1}^{n} X_i \qquad\qquad Z_n = \frac{S_n - n\mu}{\sigma\sqrt{n}}$$

where:

$$\mu \;=\; E[X_i] \qquad \sigma^2 \;=\; \mathrm{Var}[X_i] \qquad (\text{every } i)$$

So Z_n is the standardization of S_n.

Let Z be a standard normal random variable, $Z \sim \mathcal{N}(0,1)$.

Central Limit Theorem

Suppose $S_n \;=\; \sum_{i=1}^{n} X_i$ for IID variables X_i, and Z_n are the standardizations of S_n.

Then for any interval $[a,b] \subset \mathbb{R}$:

$$\lim_{n\to\infty} \; P\big[\, a \le Z_n \le b \,\big] \;=\; \Phi(b) - \Phi(a)$$

We say that Z_n *converges in probability* to the standard normal Z.

The distribution of *a very large sum* of IID variables is determined merely by μ and σ^2 from the original IID variables, while the data of higher moments fades away.

The name "**normal distribution**" is used because it arises from a large sum of repetitions of *any* other kind of distribution. It is therefore ubiquitous in applications.

Misuse of the CLT

It is important to learn when the CLT is applicable and when it is not. Many people (even professionals) apply it wrongly.

For example, sometimes one hears the claim that *if enough students take an exam, the distribution of scores will be approximately normal*. This is totally wrong!

Intuition for the CLT

The CLT is about the *distribution of simultaneity*, or (in other words) about *accumulated alignment* between independent variables.

With a large n, deviations of the total sum are predominantly created by *simultaneous* (correlated) deviations of a large portion of summands away from their means, rather than the contributions of individual summands deviating a large amount.

Simultaneity across a large n of independent items is described by... the bell curve.

04 Illustration

Exercise - Test scores distribution

Explain what is wrong with the claim that test scores should be normally distributed when a large number of students take a test.

Can you imagine a scenario with a *good* argument that test scores would be normally distributed?

(Hint: think about the composition of a single test instead of the number of students taking the test.)

Exercise - Height follows a bell curve

The height of female American basketball players follows a bell curve. Why?

05 Theory - extra

Extra - Moment Generating Functions

In order to show why the CLT is true, we introduce the technique of **moment generating functions**. Recall that the n^{th} moment of a distribution X is simply $E[X^n]$. Write μ_n for this value.

Recall the power series for e^x:

$$e^x = 1 + x + \frac{1}{2!}x^2 + \frac{1}{3!}x^3 + \dots$$

The function $f(x) = e^x$ has the property of being a bijective differentiable map from $\mathbb{R}$ to $\mathbb{R}^{>0}$, and it converts addition to multiplication: $e^{x+y} = e^x \cdot e^y$.

Given a random variable X, we can compose X with $f(x) = e^x$ to obtain a new variable. Define the **moment generating function of X** as follows:

$$M_X(t) = E[e^{tX}].$$

This is a function of $t \in \mathbb{R}$ and returns values in $\mathbb{R}$. It is called the moment generating function because it contains the data of all the higher moments μ_n. They can be extracted by taking derivatives and evaluating at zero:

$$\begin{aligned} M_X(t) &= 1 + E[X]t + E[X^2]\frac{t^2}{2!} + E[X^3]\frac{t^3}{3!} + \dots \\ M_X^{(n)}(0) &= E[X^n] = \mu_n. \end{aligned}$$

It is reasonable to consider $M_X(t)$ as a formal power series in the variable t that has the higher moments for coefficients.

Example - Moment generating function of a standard normal

We compute $M_Z(t)$ where $Z \sim \mathcal{N}(0,1)$. From the formula for expected value of a function of a random variable, we have:

$$E[e^{tZ}] = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{+\infty} e^{tx - x^2/2}\, dx.$$

Complete the square in the exponent: $tx - x^2/2 = -\frac{1}{2}(x-t)^2 + \frac{1}{2}t^2$. Thus:

$$e^{tx-x^2/2} = e^{-\frac{1}{2}(x-t)^2} e^{\frac{1}{2}t^2}.$$

The last factor can be taken outside the integral:

$$E[e^{tZ}] = e^{t^2/2} \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{+\infty} e^{-\frac{1}{2}(x-t)^2}\, dx = e^{t^2/2} = M_Z(t).$$

Exercise - Moment generating function of an exponential variable

Compute $M_X(t)$ for $X \sim \text{Exp}(\lambda)$.

Moment generating functions have the remarkable property of encoding the distribution itself:

Distributions determined by MGFs

Assume $M_X(t)$ and $M_Y(t)$ both converge. If $M_X(t) = M_Y(t)$, then $X \sim Y$.

Moreover, if $M_X(t) = M_Y(t)$ for any interval of values $t \in (-\varepsilon, \varepsilon)$, then $M_X(t) = M_Y(t)$ for all t and $X \sim Y$.

⚠ Be careful about moments vs. generating functions!

Sometimes the moments all exist, but they grow so fast that the moment generating function does not converge. For example, the log-normal distribution e^Z for $Z \sim \mathcal{N}(0,1)$ has this property.

The fact above does not apply when this happens.

When moment generating functions *approximate* each other, their corresponding distributions also approximate each other:

Distributions converge when MGFs converge

Suppose that $M_{X_n}(t) \to M_X(t)$ for all t on some interval $t \in (-\varepsilon, +\varepsilon)$. (In particular, assume that $M_X(t)$ converges on some such interval.) Then for any $[a, b]$, we have:

$$\lim_{n\to\infty} P(X_n \in [a,b]) = P(X \in [a,b]).$$

Exercise Using an MGF

Suppose X is nonnegative and $M_X(t) = (1-2t)^{-3/2}$ when $t < 1/2$ and $M_X(t) = \infty$ when $t \geq 1/2$. Find a bound on $P(X > 8)$ using (a) Markov's Inequality, and (b) Chebyshev's Inequality.

Extra - Derivation of CLT ›

The main role of moment generating functions in the proof of the CLT is to convert the sum $X_1 + \cdots + X_n$ into a product $e^{X_1} \cdots e^{X_n}$ by putting the sum into an exponent.

We have $S_n = X_1 + \cdots + X_n$, and recall $Z_n = \frac{S_n - n\mu}{\sigma\sqrt{n}}$, so $E[Z_n] = 0$ and $\mathrm{Var}(Z_n) = 1$. First, compute the MGF of Z_n. We have:

$$M_{Z_n}(t) = E[e^{tZ_n}] = E\left[e^{t\frac{S_n - n\mu}{\sigma\sqrt{n}}}\right]$$

Exchange the sum in the exponent for a product of exponentials:

$$\begin{aligned}\exp\left(t\frac{S_n - n\mu}{\sigma\sqrt{n}}\right) &= \exp\left(\frac{t}{\sigma\sqrt{n}}\sum_{i=1}^{n}(X_i - \mu)\right) \\ &= \prod_{i=1}^{n}\exp\left(\frac{t}{\sigma\sqrt{n}}(X_i - \mu)\right)\end{aligned}$$

Now since the X_i are independent, the factors $\exp\left(\frac{t}{\sigma\sqrt{n}}(X_i - \mu)\right)$ are also independent of each other. Use the product rule $E[XY] = E[X]E[Y]$ when X, Y are independent to obtain:

$$M_{Z_n}(t) = \prod_{i=1}^{n} E\left[\exp\left(\frac{t}{\sigma\sqrt{n}}(X_i - \mu)\right)\right]$$

Now expand the exponential in its Taylor series and use linearity of expectation:

$$\begin{aligned}M_{Z_n}(t) &= \prod_{i=1}^{n} E\left[1 + \frac{t}{\sigma\sqrt{n}}(X_i - \mu) + \frac{1}{2!}\left(\frac{t}{\sigma\sqrt{n}}\right)^2(X_i - \mu)^2 + \ldots\right] \\ &= \prod_{i=1}^{n}\left(1 + \frac{t}{\sigma\sqrt{n}}E[X_i - \mu] + \frac{1}{2!}\left(\frac{t}{\sigma\sqrt{n}}\right)^2 E[(X_i - \mu)^2] + \ldots\right) \\ &= \prod_{i=1}^{n}\left(1 + 0 + \frac{t^2}{2!n} + \ldots\right) \\ &\approx \prod_{i=1}^{n}\left(1 + \frac{t^2}{2n}\right).\end{aligned}$$

We don't give a complete argument for the final approximation, but a few remarks are worthwhile. For fixed n, σ, μ, and assuming the moments $E[(X_i - \mu)^k]$ have adequately bounded growth in k, the series in each factor converges for all t. Using Taylor's theorem we could write an error term as a shrinking function of n. The real trick of analysis is to show that in the product of n factors, these error terms shrink fast enough that the limit value is not affected.

In any case, the factors of the last line are independent of n, so we have:

$$M_{Z_n}(t) \approx \left(1 + \frac{t^2}{2n}\right)^n \quad \xrightarrow[n\to\infty]{} \quad \exp\left(\frac{t^2}{2}\right)$$

But $e^{\frac{t^2}{2}}$ is the MGF of $Z \sim \mathcal{N}(0, 1)$. Therefore $M_{Z_n}(t) \to M_Z(t)$, so $F_{Z_n}(x) \to F_Z(x)$.

06 Theory

Normal approximations rely on the limit stated in the CLT to approximate probabilities for large sums of variables.

Normal approximation of binomial

Let $S_n = X_1 + \cdots + X_n$ for IID variables X_i with $\mu = E[X_i]$ and $\sigma^2 = \text{Var}[X_i]$.

The **normal approximation** of S_n is:

$$F_{S_n}(s) \approx \Phi\left(\frac{s - n\mu}{\sigma\sqrt{n}}\right)$$

For example, suppose $X_i \sim \text{Ber}(p)$, so $S_n \sim \text{Bin}(n, p)$. We know $\mu = p$ and $\sigma = pq$. Therefore:

$$F_{S_n}(s) \approx \Phi\left(\frac{s - np}{\sqrt{npq}}\right)$$

A rule of thumb is that the normal approximation to the binomial is effective when $npq > 10$.

Efficient computation

This CDF is *far* easier to compute for large n than the CDF of S_n itself. The factorials in $\binom{n}{k}$ are hard (even for a computer) when n is large, and the summation adds another n factor to the scaling cost.

07 Illustration

Example - Binomial estimation: 10,000 flips

Flip a fair coin 10,000 times. Write H for the number of heads.

Estimate the probability that $4850 < H < 5100$.

Solution

(1) Check the rule of thumb: $p = q = 0.5$ and $n = 10,000$, so $npq = 2500 \gg 10$ and the approximation is effective.

---

(2) Now, calculate needed quantities:

$$\mu = E[X_i] \quad \gg\gg \quad \mu = 0.5 \quad \gg\gg \quad n\mu = 5000$$

$$\sigma^2 = \text{Var}[X_i] \quad \gg\gg \quad \sigma = 0.5 \quad \gg\gg \quad \sigma\sqrt{n} = 50$$

---

(3) Set up CDF:

$$F_H(h) = \Phi\left(\frac{h - 5000}{50}\right)$$

---

(4) Compute desired probability:

$$P[4850 < H < 5100] \quad = \quad F_H(5100) - F_H(4850)$$

$$\gg\gg \quad \Phi\left(\frac{100}{50}\right) - \Phi\left(\frac{-150}{50}\right) \quad \gg\gg \quad \Phi(2) - \Phi(-3)$$

$$\gg\gg \quad \approx 0.9772 - (1 - 0.9987) \quad \gg\gg \quad \boxed{0.9759}$$

Example - Summing 1000 dice

Suppose 1,000 dice are rolled.

Estimate the probability that the total sum of rolled numbers is more than 3,600.

Solution

(1) Let X_i be the number rolled on the i^{th} die.

Let $S = \sum_{i=1}^{n} X_i$, so S sums up the rolled numbers.

We seek $P[S \geq 3600]$.

---

(2) Now, calculate needed quantities:

$$\mu = E[X_i] \quad \gg\gg \quad \mu = 7/2 \quad \gg\gg \quad n\mu = 3500$$

$$\sigma^2 = \text{Var}[X_i] \quad \gg\gg \quad \sigma = \sqrt{\frac{35}{12}} \quad \gg\gg \quad \sigma\sqrt{n} = \sqrt{\frac{35000}{12}}$$

---

(3) Set up CDF:

$$F_S(s) \;\approx\; \Phi\left(\frac{s - 3500}{\sqrt{\frac{35000}{12}}}\right)$$

---

(4) Compute desired probability:

$$P[S \geq 3600] \;=\; 1 - F_S(3600)$$

$$\gg\gg \quad \approx 1 - \Phi\left(\frac{100}{54.01}\right) \quad \gg\gg \quad 1 - \Phi(1.852) \approx \boxed{0.968}$$

Exercise - Estimating S_{1000}

The odds of a random poker hand containing one pair is 0.42.

Estimate the probability that at least 450 out of 1000 poker hands will contain one pair.

Exercise - Nutrition study

A nutrition review board will endorse a diet if it has any positive effect in at least 65% of those tested in a certain study with 100 participants.

Suppose the diet is bogus, but 50% of participants display some positive effect by pure chance.

What is the probability that it will be endorsed?

Answer

$0.0019 \;=\; 1 - \Phi(2.9)$