Statistics and Probabilities for Computing (CEN1602)

Unit-7

Statistical Inference: Confidence Intervals and Hypothesis Testing

Prepared by

Dr. Ahmad Alzahrani

Outline

1. From Unit 6 to Unit 7

- Sampling distributions, standard error, and the Central Limit Theorem
- How sample statistics are used to make conclusions about population parameters

2. Statistical inference

- Estimation versus decision-making

3. Confidence intervals

- Confidence level and margin of error
- Confidence interval for the population mean
- Z interval when σ is known
- Finding the critical value
- t interval when σ is unknown
- Confidence interval for a population proportion

4. Hypothesis testing

- Null hypothesis and alternative hypothesis
- One-sided and two-sided tests
- Type I and Type II errors
- Significance level α
- Test statistic, rejection region, and p-value
- Tests for the population mean
- Tests for a population proportion

5. Unit summary

Learning Objectives for Unit 7

By the end of this unit, students should be able to:

- Explain the purpose of **statistical inference**.
- Distinguish between **estimation** and **hypothesis testing**.
- Explain the meaning of a **confidence interval**.
- Interpret **confidence level** and **margin of error**.
- Construct and interpret a confidence interval for a population mean.
- Decide when to use the $\boldsymbol{Z}$ **interval** and when to use the $\boldsymbol{t}$ **interval**.
- Construct and interpret a confidence interval for a population proportion.
- Define the **null hypothesis** $\boldsymbol{H_0}$ and **alternative hypothesis** $\boldsymbol{H_1}$.
- Distinguish between one-sided and two-sided hypothesis tests.
- Explain Type I error, Type II error, and **significance level** $\boldsymbol{\alpha}$.
- Compute and interpret a test statistic.
- Use the rejection-region method and the p-value method to make hypothesis-testing decisions.
- Perform basic hypothesis tests for a population mean and a population proportion.

From Unit 6 to Unit 7

- In Unit 6, we studied that if many samples of the same size n are selected from a normal population with μ and σ, each sample may give a different value of sample mean, $\bar{X}$, and thus we get a sampling distribution of $\bar{X}$.
- By the Central Limit Theorem, for a large enough sample size ($n \geq 30$), the sampling distribution of $\bar{X}$ is approximately normal with:

$$\mu_{\bar{X}} = E(\bar{X}) = \mu \qquad \sigma_{\bar{X}} = \text{SE} = \frac{\sigma}{\sqrt{n}}$$

- In Unit 7, we use sampling distribution to answer two inference questions:

How close is a sample mean $\bar{x}$ likely to be to μ?

- This leads to *Confidence Intervals*.
- A confidence interval gives a range of plausible values for the unknown population parameter μ.

Is the observed $\bar{x}$ far enough from a claimed μ_0?

- This leads to *Hypothesis Testing*.
- Hypothesis testing uses sample evidence to decide whether a claim about the population is reasonable or should be rejected.
- In short:

Sample data → Statistical inference → Conclusion about the population

Two Types of Statistical Inference

Statistical inference uses sample data to make **conclusions** about an unknown population parameter. In this unit, we focus on two main types:

1. Estimation

In Unit 6, we introduced **point estimation** which gives one numerical value to estimate an unknown parameter:

$$\bar{x} = 205\ ms$$

This means 205 ms is our point estimate of the unknown population mean.

In Unit 7, we extend this to **interval estimation**. An interval estimate gives a range of plausible values:

$$L \leq \mu \leq U$$

where L is the lower confidence limit and U is the upper confidence limit.

This range is called a **confidence interval (CI)**.

2. Hypothesis Testing

Hypothesis testing asks:

Is *there enough sample evidence against a specific claim?*

Example: $H_1: \mu < 200$

This hypothesis claims that the true mean response time is less than 200 ms. Then sample data is used to decide whether this claim is reasonable.

In short: Estimation gives a range, hypothesis testing gives a decision.

Confidence Interval, Confidence Level, and Margin of Error

- **Confidence Interval (CI)** is determined using a sample estimate as follows:

$$CI = sample\ estimate \pm margin\ of\ error$$

- For estimating a population mean:

$$CI = \bar{x} \pm E$$

where $\bar{x} = sample\ mean$ and $E = margin\ of\ error$

- The **margin of error** is the distance added and subtracted from the sample estimate. For example:

$$205 \pm 10$$

- gives:

$$195 \leq \mu \leq 215$$

- A CI is reported with a **confidence level**, such as:

$$90\%\ \text{CI},\ 95\%\ \text{CI},\ 99\%\ \text{CI}$$

- The **confidence level** describes the reliability of the CI calculation. A 95% confidence level means that if we repeatedly take samples and calculate a CI each time, about 95% of those intervals will contain the true parameter.
- In general:

higher confidence level → larger margin of error

larger margin of error → wider CI

CI for the Population Mean When σ is Known

- When estimating the population mean μ, and the population standard deviation σ is known, we use the **Z confidence interval** :

$$\boldsymbol{CI = \bar{x} \pm z_{\alpha/2} \times \mathrm{SE}}$$

- where: $\bar{x} =$ *sample mean*, $\sigma =$ *population standar deviation*

$n =$ *sample size*, *standard error* (SE)= $\frac{\sigma}{\sqrt{n}}$

$z_{\alpha/2}$= the **critical value**

- The margin of error here is $E = z_{\alpha/2} \times \mathrm{SE}$
- The **critical value** is chosen based on the confidence level:

$$90\% \rightarrow z_{\alpha/2} = 1.645$$
$$95\% \rightarrow z_{\alpha/2} = 1.96$$
$$99\% \rightarrow z_{\alpha/2} = 2.576$$

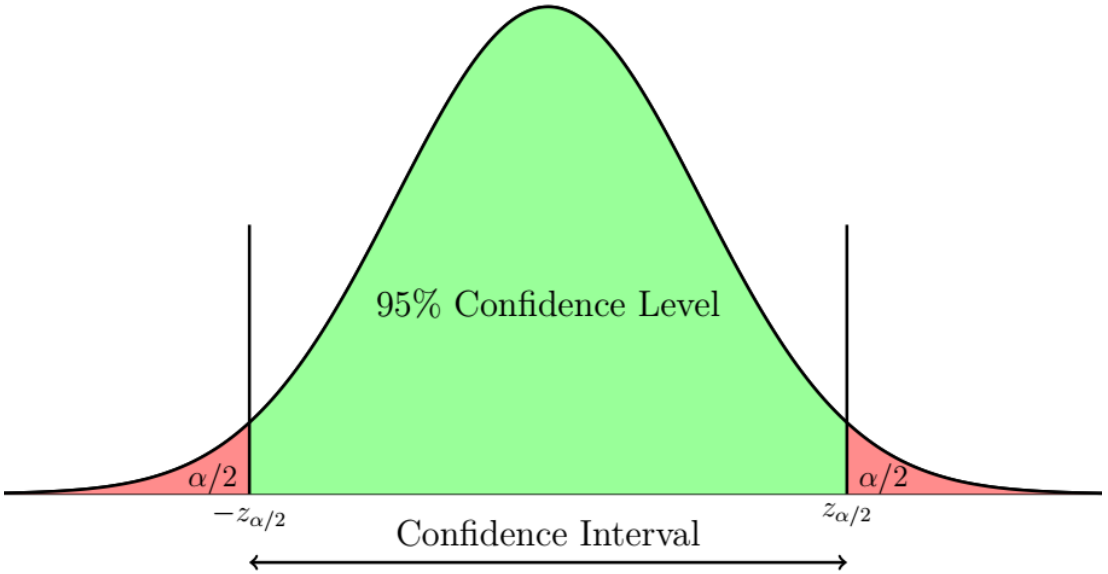
- A higher confidence level uses a larger critical value, so the interval becomes wider.

Finding the Critical Value

- The critical value $z_{\alpha/2}$ is found from the confidence level. In general:

$$\alpha = 1 - \textit{confidence level}$$

- α is the total area left outside the confidence interval, and it is often is called the **significance level.**



- Since the confidence interval is two-sided, $\frac{\alpha}{2}$ is placed in each tail. The critical value $z_{\alpha/2}$ is the z-value that leaves $\alpha/2$ in the upper tail.
- **For example, for a 95% confidence interval:**

$$1 - \alpha = 0.95 \;\rightarrow\; \alpha = 0.05 \;\rightarrow\; \frac{\alpha}{2} = 0.025$$

$$z_{\alpha/2} = z_{0.025} = 1.96 \text{ (from the probability table)}$$

Example: Z Confidence Interval for the Mean

A web server is tested during peak hours. A random sample of $n = 36$ requests has a sample mean response time of :

$$\bar{x} = 205\ ms$$

Assume the population standard deviation is known:

$$\sigma = 30\ ms$$

Construct a **95%** confidence interval for the true mean response time μ.

- For 95% confidence, $z_{\alpha/2} = 1.96$
- Use:

$$\bar{x} \pm z_{\alpha/2}\text{SE} = \bar{x} \pm z_{\alpha/2}\frac{\sigma}{\sqrt{n}}$$

$$205 \pm 1.96\frac{30}{\sqrt{36}}$$

$$205 \pm 1.96(5)$$

$$205 \pm 9.8$$

So:

$$195.2 \leq \mu \leq 214.8$$

- Therefore, the 95% CI is:

$$[195.2, 214.8]\ ms$$

Interpreting the Z Confidence Interval

- From the previous example, the 95% CI was:

$$195.2 \leq \mu \leq 214.8$$

- This interval estimates:

$$\mu = \textit{true mean response time of all web server requests}$$

- A simple interpretation is:
  We are 95% confident that the true mean response time is between 195.2 ms and 214.8 ms.
- The interval is centered around the sample mean:

$$\bar{x} = 205\ ms$$

- The margin of error is:

$$E = 9.8\ ms$$

- This means the interval extends 9.8 ms below and above the sample mean.
- If the target mean response time is 200 ms, then 200 lies inside the interval.
- So, based on this interval alone, the sample does not clearly show that the true mean is above 200 ms.

CI for the Population Mean When σ is Unknown

- In the previous confidence interval, we assumed that the population standard deviation σ was known.
- In practice, σ is often unknown.
- Instead, we calculate the sample standard deviation s from the sample data.
- So, instead of using $\frac{\sigma}{\sqrt{n}}$
- we use $\frac{s}{\sqrt{n}}$
- This changes the confidence interval method.
- When σ is unknown, we usually use the ***t* distribution** instead of the standard normal distribution.
- The next confidence interval is called the:

  t confidence interval for the population mean

- This is the most common case in real data analysis because population standard deviation is rarely known.

The t Distribution

- The t distribution is similar to the standard normal distribution, but it has wider tails.
- Wider tails mean it allows more uncertainty when the sample size is small.
- The shape of the t distribution depends on the degrees of freedom:

$$df = n - 1$$

- Examples:

$$n = 10 \Rightarrow df = 9$$
$$n = 25 \Rightarrow df = 24$$

- As n becomes larger, the t distribution becomes closer to the standard normal distribution.
- So, when σ is unknown, we use:

$$t_{\alpha/2,\ n-1}$$

- instead of:

$$z_{\alpha/2}$$

CI for the Population Mean When σ is Unknown

- When estimating the population mean μ, and the population standard deviation σ is unknown, we use the **t confidence interval**:

$$\bar{x} \pm t_{\alpha/2,\ n-1} \times SE$$

where:

$\bar{x} =$ *sample mean* $\quad s =$ *sample standard deviation*

$n =$ *sample size* $\quad df = n - 1$

$t_{\alpha/2,\ n-1} =$ *critical value from the t distribution*

- The standard error is:

$$SE = \frac{s}{\sqrt{n}}$$

- The margin of error is:

$$E = t_{\alpha/2,\ n-1} \frac{s}{\sqrt{n}}$$

- So the interval can also be written as:

$$\bar{x} \pm E$$

- This interval is used more often than the Z interval because σ is usually unknown in real data.

Common t Critical Values for Confidence Intervals

Sample Size (n)	Degrees of Freedom ($df = n - 1$)	90% CI	95% CI	99% CI
3	2	2.920	4.303	9.925
4	3	2.353	3.182	5.841
5	4	2.132	2.776	4.604
6	5	2.015	2.571	4.032
7	6	1.943	2.447	3.707
8	7	1.895	2.365	3.499
9	8	1.860	2.306	3.355
10	9	1.833	2.262	3.250
11	10	1.812	2.228	3.169
12	11	1.796	2.201	3.106
13	12	1.782	2.179	3.055
14	13	1.771	2.160	3.012
15	14	1.761	2.145	2.977
16	15	1.753	2.131	2.947
20	19	1.729	2.093	2.861
25	24	1.711	2.064	2.797
30	29	1.699	2.045	2.756
40	39	1.685	2.023	2.708
50	49	1.677	2.010	2.680
100	99	1.660	1.984	2.626
Large (n)	∞	**1.645**	**1.960**	**2.576**

Notice: as n becomes larger, the t critical value gets closer to the Z critical value.

Example: t Confidence Interval for the Mean

A web server is tested during peak hours. A random sample of $n = 16$ requests has $\bar{x} = 205$ *ms*, and sample standard deviation $s = 24$ *ms*.
Construct a **95% confidence interval** for the true mean response time μ.

- Since σ is unknown, use the t confidence interval:

$$\bar{x} \pm t_{\alpha/2,\ n-1} \frac{s}{\sqrt{n}}$$

- Degrees of freedom $df = n - 1 = 16 - 1 = 15$
- For 95% confidence and $df = 15$:

$$t_{\alpha/2,\ 15} = t_{0.025,\ 15} = 2.131$$

- Now calculate:

$$205 \pm 2.131 \frac{24}{\sqrt{16}}$$
$$205 \pm 2.131(6)$$
$$205 \pm 12.786$$

- So $192.214 \leq \mu \leq 217.786$
- Therefore, the 95% t confidence interval is approximately:

$$[192.2, 217.8]\ ms$$

Confidence Interval for a Population Proportion

- Sometimes the parameter of interest is a population proportion:

$$p$$

- A proportion is used when the outcome has two categories, such as:

failed request / successful request

vulnerable device / secure device

- From the sample, we calculate the sample proportion:

$$\hat{p} = \frac{x}{n}$$

- where: $x = number\ of\ successes\ in\ the\ sample$

$n = sample\ size$

- Similar to the confidence interval for the mean, we use:

$$CI = \text{ sample estimate} \pm \text{margin of error}$$

- For a population proportion, the confidence interval is:

$$CI = \hat{p} \pm z_{\alpha/2}\sqrt{\frac{\hat{p}\hat{q}}{n}}$$

- Here $\hat{p} = sample\ estimate\ of\ p,$ and $\hat{q} = 1 - \hat{p}$
- So, the idea is the same: use a sample statistic to estimate an unknown population parameter.

Example: Confidence Interval for a Population Proportion

A system administrator checks a random sample of $n = 200$ web requests. Suppose $x = 30$ requests were slow. The sample proportion is:

$$\hat{p} = \frac{x}{n} = \frac{30}{200} = 0.15$$

So, 15% of the sampled requests were slow. Construct a **95% confidence interval** for the true proportion of slow requests.

For 95% confidence $z_{\alpha/2} = 1.96$

Use:

$$\hat{p} \pm z_{\alpha/2}\sqrt{\frac{\hat{p}\hat{q}}{n}}$$

$$0.15 \pm 1.96\sqrt{\frac{0.15(1-0.15)}{200}}$$

$$0.15 \pm 0.0495$$

So $0.1005 \leq p \leq 0.1995$

Therefore, the 95% CI is approximately:

$$10.1\% \leq p \leq 20.0\%$$

This estimates the true proportion of slow requests in the system.

Confidence Intervals: Summary

- A confidence interval estimates an unknown population parameter using sample data.
- In this unit, we covered three common cases:

- ***Mean with known σ:***

$$\bar{x} \pm z_{\alpha/2} \frac{\sigma}{\sqrt{n}}$$

- ***Mean with unknown σ:***

$$\bar{x} \pm t_{\alpha/2,n-1} \frac{s}{\sqrt{n}}$$

- ***Population proportion***

$$\hat{p} \pm z_{\alpha/2} \sqrt{\frac{\hat{p}\hat{q}}{n}}$$

- The common idea is:

$$CI = \text{sample estimate} \pm \text{margin of error}$$

Hypothesis Testing

- **Hypothesis testing** is used to check a **claim** about a population parameter.
- Instead of estimating the parameter with an interval, we start with a specific claim and test whether the sample data supports or contradicts it.
- For example, suppose the maximum acceptable true mean response time for a web server is:

$$\mu = 200\ ms$$

- A hypothesis test asks:

  Does *the sample provide enough evidence that the true mean response time is above* 200 *ms?*

- The decision is based on how far the sample result is from the claimed value.
- In hypothesis testing, sample data is treated as evidence.

  claim *about population* → *sample evidence* → *statistical decision* .

Null and Alternative Hypotheses

- A hypothesis test usually has **two competing statements**:

$$H_0 = \textit{null hypothesis}$$
$$H_1 = \textit{alternative hypothesis}$$

- The **null hypothesis** represents the starting claim or default assumption.
- The **alternative hypothesis** represents what we are trying to find evidence for.
- For example, suppose the maximum acceptable true mean response time for a web server is:

$$200\ ms$$

- We want to test whether the server is slower than this limit.

$$H_0: \mu \leq 200$$
$$H_1: \mu > 200$$

- H_0 means the true mean response time is within the acceptable limit.
- H_1 means the true mean response time is above the acceptable limit, meaning the server may be slower than expected.
- In hypothesis testing, the sample data is used to decide whether there is enough evidence to reject H_0.

One-Sided and Two-Sided Tests

The form of the alternative hypothesis H_1 determines the type of test.

1. Right-sided test

Used when we want to check whether the parameter is **greater than** a reference value ($H_1: \mu > \mu_0$).

Example:

$$H_1: \mu > 200$$

This checks whether the true mean response time is above the acceptable limit.

2. Left-sided test

Used when we want to check whether the parameter is **less than** a reference value ($H_1: \mu < \mu_0$).

Example:

$$H_1: \mu < 200$$

This checks whether the true mean response time is below a reference value.

3. Two-sided test

Used when we want to check whether the parameter is **different from** a reference value ($H_1: \mu \neq \mu_0$)

Example:

$$H_1: \mu \neq 200$$

This checks whether the true mean response time is different from 200ms in either direction.

Type I and Type II Errors

In hypothesis testing, the decision may be correct or incorrect. There are two possible errors:

1. Type I Error

A Type I error happens when we reject H_0 even though H_0 is actually true.

$$\text{Type } I \text{ error} = \text{rejecting a true } H_0$$

In the web server example:

$$H_0: \mu \leq 200$$

A Type I error means we conclude that the server is too slow, even though its true mean response time is within the acceptable limit.

2. Type II Error

A Type II error happens when we fail to reject H_0 even though H_0 is actually false.

$$\text{Type } II \text{ error} = \text{not rejecting a false } H_0$$

In the web server example, this means we fail to detect that the server is too slow.

- We use Alpha (α) and beta (β) as the probabilities of making Type I and Type II errors, respectively:

$$\alpha = P(\text{Type I error})$$
$$\beta = P(\text{Type II error})$$

Significance Level α in Hypothesis Testing

- Earlier, in confidence intervals, we used $\boldsymbol{\alpha}$ as a significance level as follows:

$$\boldsymbol{\alpha = 1 - \textit{confidence level}}$$

- For example, for a 95% confidence interval:

$$\alpha = 1 - 0.95 = 0.05$$

- In hypothesis testing, α is also used, but now it represents the **maximum probability allowed for making a Type I error**.

$$\alpha = P(\textit{reject } H_0 \textit{ when } H_0 \textit{ is true})$$

- Common confidence levels and significance levels are:

$$90\% \text{ confidence level} \rightarrow \alpha = 0.10$$
$$95\% \textit{ confidence level} \rightarrow \alpha = 0.05$$
$$99\% \textit{ confidence level} \rightarrow \alpha = 0.01$$

- For example:

$$\alpha = 0.05$$

  means the test allows a maximum 5% probability of wrongly rejecting a true H_0.

- A smaller α means stronger evidence is needed to reject H_0.

Test Statistic

- A **test statistic** converts the sample result into a standardized value.
- It measures how far the sample result is from the value stated in H_0.
- For the cases covered in this unit, the general idea is:

$$\textbf{test } \textit{statistic} = \frac{\textit{sample estimate} - \textit{claimed value}}{\textit{standard error}}$$

Test case	Sample estimate	Claimed value	Standard error	Test statistic
Population mean σ known	$\bar{x}$	μ_0	$\frac{\sigma}{\sqrt{n}}$	$\boldsymbol{z = \frac{\bar{x} - \mu_0}{\sigma/\sqrt{n}}}$
Population proportion	$\hat{p}$	p_0	$\sqrt{\frac{p_0 q_0}{n}}$	$\boldsymbol{z = \frac{\hat{p} - p_0}{\sqrt{\frac{p_0 q_0}{n}}}}$

- A large positive or negative test statistic means the sample result is far from the claim.
- So, the test statistic tells us how unusual the sample result is if H_0 is true.

Hypothesis Test Decision: Rejection Region Method

- After computing the test statistic, we compare it with a critical value.
- The **rejection region** is the area where the test statistic is too far from what H_0 expects.
- For a right-sided Z-test:

$$H_1: \mu > \mu_0$$

- we reject H_0 if:

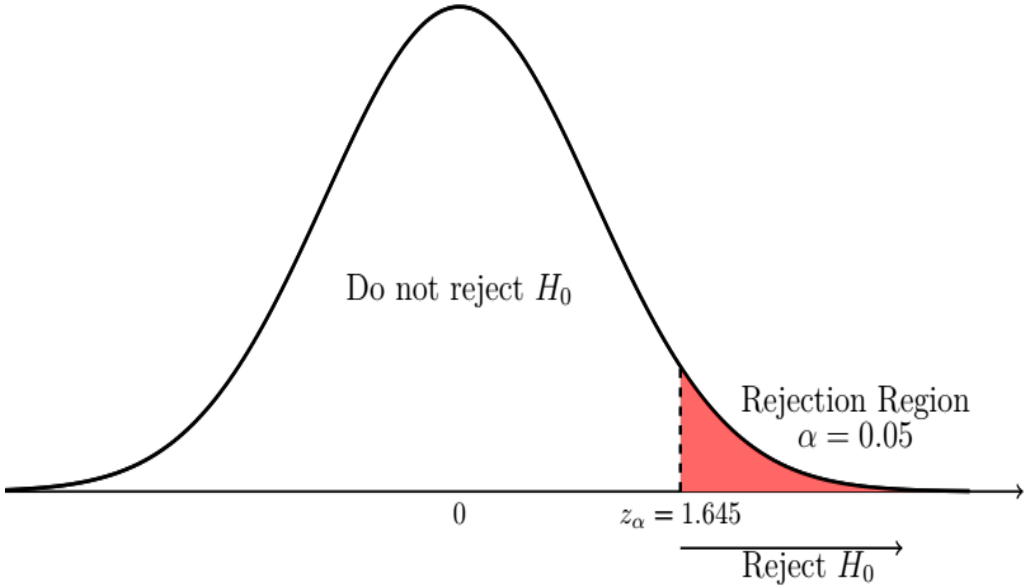
$$z > z_\alpha$$

- For example, if:

$$\alpha = 0.05$$

- then:

$$z_\alpha = 1.645$$



- So, the decision rule is:

Reject H_0 *if* $z > 1.645$

- If the test statistic falls in the rejection region, the sample gives strong evidence against H_0.
- If it does not fall in the rejection region, we fail to reject H_0.
- **Note:** For this right-sided test, all $\alpha = 0.05$ is in one tail, so $z_\alpha = 1.645$ .In a two-sided CI as in slide 8, α is split into two tails, so we use $z_{\alpha/2} = 1.96$.

Rejection Region Method: Example 1

A web server should have a mean response time not greater than 200 *ms*. A sample of $n = 36$ requests gives:

$$\bar{x} = 212 \text{ ms}, \qquad \sigma = 30 \text{ ms}$$

At the 95% confidence level, test whether the true mean response time is greater than 200 ms.

Hypotheses:

$$H_0: \mu \leq 200$$

$$H_1: \mu > 200$$

$$\alpha = 1 - 0.95 = 0.05$$

Test statistic:

$$z = \frac{\bar{x} - \mu_0}{\sigma/\sqrt{n}} = \frac{212 - 200}{30/\sqrt{36}} = \frac{12}{5} = 2.40$$

For a right-sided test with $\alpha = 0.05$:

$$z_\alpha = 1.645$$

Since $2.40 > 1.645$, we reject H_0.

Conclusion: There is enough evidence that the true mean response time is greater than 200ms.

Rejection Region Method: Example 2

A chip manufacturer requires the defect rate to be $\leq 5\%$. A quality-control engineer tests a random sample of $n = 200$ chips and finds $x = 16$ defective chips.
The sample defect proportion is:

$$\hat{p} = \frac{16}{200} = 0.08$$

At the **95% confidence level**, test whether the true defect rate is greater than 5%.
Convert confidence level to significance level:

$$\alpha = 1 - 0.95 = 0.05$$

Hypotheses: $H_0: p \leq 0.05$
$H_1: p > 0.05$

Since this is a right-sided test, the rejection rule is:

Reject H_0 *if* $z > z_\alpha$

For: $\alpha = 0.05,\ z_\alpha = 1.645$
Test statistic:

$$z = \frac{\hat{p} - p_0}{\sqrt{\frac{p_0 q_0}{n}}} = \frac{0.08 - 0.05}{\sqrt{\frac{0.05(0.95)}{200}}} = \frac{0.03}{0.0154} = 1.95$$

Decision $1.95 > 1.645$, so Reject H_0.

Conclusion: There is enough evidence that the true chip defect rate is greater than 5%. This suggests that the manufacturing process may not be meeting the required quality level.

Hypothesis Test Decision: p-value Method

- The ***p*-value** measures how strong the sample evidence is against H_0.
- It answers this question:

  If H_0 *is true, how unusual is the sample result?*
- A small p-value means the sample result is unlikely under H_0.
- So, a small p-value gives stronger evidence against H_0.
- Decision rule:

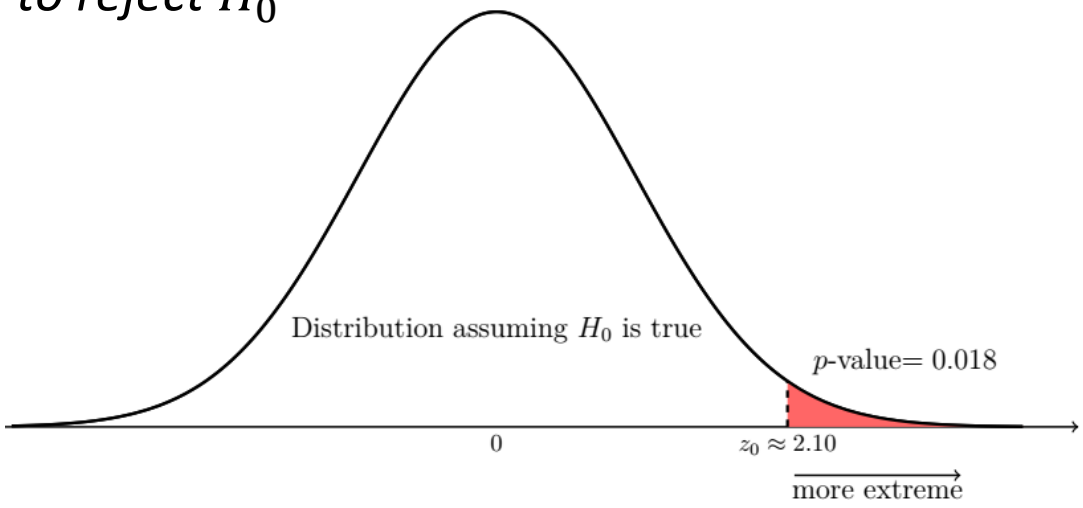
  If $p\text{–}value \leq \alpha$, *reject* H_0

  If $p\text{–}value > \alpha$, *fail to reject* H_0
- For example, if:

  $\alpha = 0.05$
- and:

  p–value $= 0.018$
- then:

  $0.018 < 0.05$
- So, we reject H_0.



p-value Method: Example 1

A chip design team expects the mean power consumption of a processor chip to be ≤ 5 *W*. After a design modification, the team wants to check whether the chip consumes more power than this limit. A random sample of $n = 64$ chips gives a sample mean power consumption of $\bar{x} = 5.3$ *W*. Assume the population standard deviation is known $\sigma = 0.6$ *W*.

At the **95% confidence level**, test whether the true mean power consumption is greater than 5W.

Convert confidence level to significance level:

$$\alpha = 1 - 0.95 = 0.05$$

Hypotheses: $H_0: \mu \leq 5$

$H_1: \mu > 5$

Test statistic:

$$z = \frac{\bar{x} - \mu_0}{\sigma/\sqrt{n}} = \frac{5.3 - 5}{0.6/\sqrt{64}} = \frac{0.3}{0.075} = 4.00$$

For a right-sided test:

$$p\text{–}value = P(Z \geq 4.00) \approx 0.00003$$

Since $0.00003 < 0.05$, we reject H_0.

Conclusion: There is enough evidence that the true mean power consumption of the chip is greater than 5W. The design modification may have increased power consumption and should be investigated.

p-value Method: Example 2

A mobile app is expected to consume on average less than 8% battery per hour during normal use. After a new update, the development team wants to check whether the app now consumes more battery. A random sample of $n = 100$ phones gives a sample mean battery consumption of $\bar{x} = 8.3\%$. Assume the population standard deviation is known $\sigma = 3\%$.

At the 95% confidence level, test whether the app consumes more than 8% battery per hour after the update.

Hypotheses: $H_0: \mu \leq 8$

$H_1: \mu > 8$

$$\alpha = 1 - 0.95 = 0.05$$

Test statistic:

$$z = \frac{\bar{x} - \mu_0}{\sigma/\sqrt{n}} = \frac{8.3 - 8}{3/\sqrt{100}} = \frac{0.3}{0.3} = 1.00$$

For a right-sided test:

$$p\text{–}value = P(Z \geq 1.00) \approx 0.1587$$

Since $0.1587 > 0.05$, we fail to reject H_0.

Conclusion: There is not enough evidence that the app consumes more than 8% battery per hour after the update.

Unit 7 Summary

- In this unit, we studied **statistical inference**, which uses sample data to make conclusions about unknown population parameters.
- We focused on two main tools:

  Confidence *Intervals*

  Hypothesis *Testing*
- For confidence intervals, we used the general idea:

  sample *estimate* $\pm$ *margin of error*
- We constructed confidence intervals for:

  μ *when* σ *is known*

  μ *when* σ *is unknown*

  p *for a population proportion*
- For hypothesis testing, we learned the meaning of H_0 and H_1, the possible errors, and how to make a test decision using:

  α, *test statistic*, *rejection region*, $p-$*value*
- The main idea of the unit is:

  sample *data* $\rightarrow$ *statistical evidence* $\rightarrow$ *conclusion about the population*
- These methods help us make data-based decisions in computing problems such as system performance, chip quality, power consumption, and software behavior as given in examples.