Hypothesis Testing in R: Exercises

Statistics and Big Data

Niccolò Salvini, PhD

UCSC

Academic Year 2025-2026

Course: Statistics and Big Data

Overview

Key Concepts Review: The Core Logic

Let's quickly review the main ideas of hypothesis testing.

The Core Logic

1. State a **Null Hypothesis (H_0)**, the "no effect" scenario.
2. State an **Alternative Hypothesis (H_1)**, what we want to prove.
3. Collect data and calculate a **test statistic**.
4. Calculate the **p-value**: the probability of observing our data (or more extreme) if H_0 is true.
5. If p-value $\leq \alpha$, we reject H_0.

Key Concepts Review: Types of Tests

Types of Tests

- **One-Sample Test:** Compare one group's mean to a known value.
- **Two-Sample Test:** Compare the means of two different groups.
- **ANOVA:** Compare the means of **more than two** groups.

Alternative Hypothesis (H_1)

The alternative hypothesis determines the type of test:

- $H_1 : \mu \neq \mu_0$ (Two-tailed)
- $H_1 : \mu > \mu_0$ (Right-tailed)
- $H_1 : \mu < \mu_0$ (Left-tailed)

The t.test() Command: One-Sample Test

The ‘t.test()‘ function is your main tool for comparing a sample mean to a known value when σ is unknown.

One-Sample t-test

Are our sample's grades different from a target of 25?

```
# H0: mu = 25
# H1: mu != 25
t.test(grades, mu = 25, alternative = "two.sided")
```

The t.test() Command: One-Sample Test

The ‘t.test()‘ function is your main tool for comparing a sample mean to a known value when σ is unknown.

One-Sample t-test

Are our sample's grades different from a target of 25?

```
# H0: mu = 25
# H1: mu != 25
t.test(grades, mu = 25, alternative = "two.sided")
```

Key Arguments

x: Your vector of data (e.g., ‘grades‘).
mu: The value from H_0 for a one-sample test.

The t.test() Command: Two-Sample Test

To compare the means of two different groups, we use the formula syntax.

Two-Sample t-test

Is there a difference in grades between Group A and Group B?

```
# H0: mu_A = mu_B
# H1: mu_A != mu_B
# Formula syntax: dependent_var ~ grouping_var
t.test(grades ~ group, data = my_data)
```

The t.test() Command: Two-Sample Test

To compare the means of two different groups, we use the formula syntax.

Two-Sample t-test

Is there a difference in grades between Group A and Group B?

```
# H0: mu_A = mu_B
# H1: mu_A != mu_B
# Formula syntax: dependent_var ~ grouping_var
t.test(grades ~ group, data = my_data)
```

Common Arguments

`alternative`: Specifies H_1 ("two.sided", "less", "greater").

`conf.level`: The confidence level, typically 0.95 for $\alpha = 0.05$.

For Comparing More Than Two Group Means

To test if the means of multiple groups are equal, we use Analysis of Variance (ANOVA).

For Comparing More Than Two Group Means

To test if the means of multiple groups are equal, we use Analysis of Variance (ANOVA).

Why not multiple t-tests?

Running multiple t-tests would inflate the Type I error rate (the chance of finding a difference that isn't really there). ANOVA avoids this.

For Comparing More Than Two Group Means

To test if the means of multiple groups are equal, we use Analysis of Variance (ANOVA).

Why not multiple t-tests?

Running multiple t-tests would inflate the Type I error rate (the chance of finding a difference that isn't really there). ANOVA avoids this.

ANOVA Hypotheses

- $H_0 : \mu_1 = \mu_2 = \mu_3 = \cdots = \mu_k$
- H_1: At least one mean is different.

For Comparing More Than Two Group Means

R Implementation

We use ‘aov()‘ to build the model and ‘summary()‘ to see the results.

```
# Compare grades across three different teaching methods
# H0: mu_method1 = mu_method2 = mu_method3
anova_model <- aov(grades ~ method, data = school_data)
summary(anova_model) # This shows the p-value
```

For Comparing More Than Two Group Means

R Implementation

We use ‘aov()‘ to build the model and ‘summary()‘ to see the results.

```
# Compare grades across three different teaching methods
# H0: mu_method1 = mu_method2 = mu_method3
anova_model <- aov(grades ~ method, data = school_data)
summary(anova_model) # This shows the p-value
```

Interpreting the Output

Look for the **F-statistic** and its corresponding p-value, usually labeled `Pr(>F)`. If this p-value is less than α, you reject H_0 and conclude that not all group means are equal.

Coffee Bag Weights

Scenario

A coffee company claims that its bags are filled with 500g of coffee on average. A quality inspector randomly samples 12 bags and records their weights:

`495, 503, 498, 490, 505, 492, 501, 494, 506, 496, 499, 502`

Using a significance level of $\alpha = 0.05$, is there evidence to suggest the mean weight is **different** from 500g?

Coffee Bag Weights

Scenario

A coffee company claims that its bags are filled with 500g of coffee on average. A quality inspector randomly samples 12 bags and records their weights:

`495, 503, 498, 490, 505, 492, 501, 494, 506, 496, 499, 502`

Using a significance level of $\alpha = 0.05$, is there evidence to suggest the mean weight is **different** from 500g?

Tasks

1. State the null (H_0) and alternative (H_1) hypotheses.
2. Create a vector in R with the sample weights.
3. Use the ‘t.test()‘ function to perform the hypothesis test.
4. Interpret the p-value and state your conclusion. Is the company’s claim supported?

New vs. Old Drug

Scenario

A pharmaceutical company develops a new drug to reduce recovery time from an illness. To test its effectiveness, they randomly assign patients to two groups: one receiving the new drug, the other a placebo. The recovery times (in days) are recorded.

New Drug Group: `5, 6, 6, 7, 5, 8, 6`

Placebo Group: `7, 8, 7, 9, 10, 8, 9`

Test at $\alpha = 0.05$ whether the new drug **reduces** recovery time.

New vs. Old Drug

Scenario

A pharmaceutical company develops a new drug to reduce recovery time from an illness. To test its effectiveness, they randomly assign patients to two groups: one receiving the new drug, the other a placebo. The recovery times (in days) are recorded.

New Drug Group: `5, 6, 6, 7, 5, 8, 6`

Placebo Group: `7, 8, 7, 9, 10, 8, 9`

Test at $\alpha = 0.05$ whether the new drug **reduces** recovery time.

Tasks

1. State the null (H_0) and alternative (H_1) hypotheses. (Hint: this is a one-tailed test!).
2. Create two separate vectors in R for the data.
3. Use ‘t.test()‘ to compare the two groups. What is the correct ‘alternative‘ argument?
4. Based on the p-value, what do you conclude about the drug’s effectiveness?

Comparing Teaching Methods

Scenario

A researcher wants to compare the effectiveness of three different teaching methods (A, B, C) on student exam scores. Students are randomly assigned to one of the three methods and their final exam scores are recorded.

Method A: `85, 90, 88, 82, 86`
Method B: `92, 95, 89, 93, 91`
Method C: `78, 81, 80, 84, 79`

Is there a significant difference in mean exam scores among the three teaching methods at $\alpha = 0.05$?

Comparing Teaching Methods

Scenario

A researcher wants to compare the effectiveness of three different teaching methods (A, B, C) on student exam scores. Students are randomly assigned to one of the three methods and their final exam scores are recorded.

Method A: `85, 90, 88, 82, 86`
Method B: `92, 95, 89, 93, 91`
Method C: `78, 81, 80, 84, 79`

Is there a significant difference in mean exam scores among the three teaching methods at $\alpha = 0.05$?

Tasks

1. State the null (H_0) and alternative (H_1) hypotheses for an ANOVA test.
2. Create a data frame in R with two columns: one for ‘scores‘ and one for the ‘method‘ (A, B, or C).
3. Use the ‘aov()‘ function and then ‘summary()‘ to analyze the data.

Exercise 4: Highway Patrol

A highway patrol samples 64 vehicles at a specific location, finding a mean speed of 66.2 mph with a standard deviation of 4.2 mph. The posted speed limit is 65 mph.

- Use $\alpha = 0.05$ to test the hypothesis that the mean speed is greater than 65 mph ($H_0 : \mu \leq 65$).
- Should the patrol place a radar trap at this location?

Homework: T-Tests

Exercise 4: Highway Patrol

A highway patrol samples 64 vehicles at a specific location, finding a mean speed of 66.2 mph with a standard deviation of 4.2 mph. The posted speed limit is 65 mph.

- Use $\alpha = 0.05$ to test the hypothesis that the mean speed is greater than 65 mph ($H_0 : \mu \leq 65$).
- Should the patrol place a radar trap at this location?

Exercise 5: Prawn Growth Rate

An experiment investigates the difference in growth rate of prawns fed either an artificial or natural diet.

- Using a dataset with columns for ‘growth_rate‘ and ‘diet‘, state the hypotheses for a two-sample t-test.
- Perform the test in R assuming equal variances. Do you reject the null hypothesis?

Exercise 6: First Coach

An athletics school evaluates a new training program. The times (in seconds) of 10 athletes are recorded before and after the training.

- `before = c(12.9, 13.5, 12.8, 15.6, 17.2, 19.2, 12.6, 15.3, 14.4, 11.3)`
- `after = c(12.7, 13.6, 12.0, 15.2, 16.8, 20.0, 12.0, 15.9, 16.0, 11.1)`
- Conduct a paired t-test to determine if there has been a statistically significant change in performance at a 95% confidence level.

Athletics Training

Exercise 6: First Coach

An athletics school evaluates a new training program. The times (in seconds) of 10 athletes are recorded before and after the training.

- `before = c(12.9, 13.5, 12.8, 15.6, 17.2, 19.2, 12.6, 15.3, 14.4, 11.3)`
- `after = c(12.7, 13.6, 12.0, 15.2, 16.8, 20.0, 12.0, 15.9, 16.0, 11.1)`
- Conduct a paired t-test to determine if there has been a statistically significant change in performance at a 95% confidence level.

Exercise 7: Second Coach

The first coach is fired. A new coach achieves the following results with the same athletes:

- `after = c(12.0, 12.2, 11.2, 13.0, 15.0, 15.8, 12.2, 13.4, 12.9, 11.0)`
- Using the same 'before' data, test again if the change is significant.

Homework: ANOVA

Exercise 8: Weight Loss Programs

90 people are randomly assigned to one of three weight-loss programs (A, B, or C). Their weight loss is recorded.

- Given a dataset with ‘weight_loss‘ and ‘program‘, state the hypotheses for an ANOVA test.
- Plot the data using a boxplot (‘boxplot(weight_loss program, ...)‘).
- Fit a one-way ANOVA model to test for a difference in mean weight loss among the programs.

Homework: ANOVA

Exercise 8: Weight Loss Programs

90 people are randomly assigned to one of three weight-loss programs (A, B, or C). Their weight loss is recorded.

- Given a dataset with ‘weight_loss‘ and ‘program‘, state the hypotheses for an ANOVA test.
- Plot the data using a boxplot (‘boxplot(weight_loss   program, ...)‘).
- Fit a one-way ANOVA model to test for a difference in mean weight loss among the programs.

Exercise 9: Insect Sprays

The built-in R dataset ‘InsectSprays‘ contains counts of insects after applying one of six different sprays.

- Use ANOVA to test if there is a significant difference in the mean number of insects among the different sprays.