

Learning Dunnett's Test: A Post-Hoc Analysis in R for Comparing to a Control Group

Authored by
Mohammed Iooti

November 7, 2025

RECOMMENDED CITATION

Mohammed Iooti (2025). *Learning Dunnett's Test: A Post-Hoc Analysis in R for Comparing to a Control Group*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=12187>

When conducting complex statistical analyses, particularly those involving comparisons among multiple group means, researchers often rely on the [ANOVA](#) (Analysis of Variance) framework. However, a significant result from an ANOVA only indicates that at least two groups differ; it does not specify which pairs are responsible for that difference. This necessitates a subsequent procedure known as a [post-hoc test](#).

A post-hoc test is a statistical method performed after rejecting the null hypothesis in an omnibus test, such as ANOVA, to pinpoint the specific group comparisons that are [statistically significantly different](#) from one another. While several post-hoc procedures exist, the choice depends heavily on the experimental design and the nature of the comparisons required.

If your experimental design involves several treatment groups that must all be compared against a single, defined [control group](#), the appropriate choice is [Dunnett's test](#). This powerful test is specifically designed to maintain the family-wise error rate when multiple comparisons are made solely against a control, making it superior to methods like Tukey's HSD for this specific scenario. This comprehensive tutorial details the steps required to execute and interpret Dunnett's test using the statistical programming language [R](#).

Case Study: Evaluating New Study Techniques in R

To illustrate the application of Dunnett's test, consider a scenario where a high school teacher is investigating whether two novel studying techniques (Treatment 1 and Treatment 2) can effectively improve student exam performance compared to the standard teaching method. The standard method serves as our control group.

The teacher recruits 30 students and randomly assigns them into three distinct groups, ensuring equal sample sizes for each condition. This experimental design is ideal for a one-way ANOVA followed by Dunnett's test, as the primary interest lies in comparing the effectiveness of each new technique exclusively against the control, rather than comparing the two new techniques against each other.

The structure of the study is as follows:

Control Group: 10 students follow the standard curriculum and study habits.

New Study Technique 1: 10 students employ the first novel technique.

New Study Technique 2: 10 students employ the second novel technique.

After one week of implementation, all 30 students take the exact same standardized examination. We will now proceed through the necessary steps in R--creating the dataset, visualizing the data, performing the preliminary one-way ANOVA, and finally executing Dunnett's test--to determine if either new technique yields a significantly different mean exam score compared to the control

group.

Step 1: Data Preparation and Initial Inspection

The first critical step in any statistical analysis is preparing the data structure. In R, we will create a data frame containing two variables: `technique` (the categorical predictor) and `score` (the numeric response variable).

The following R code establishes the dataset, combining the exam scores collected from the 30 students across the three study conditions. We utilize the `rep()` function to efficiently generate the factor levels for the `technique` variable, ensuring 10 observations for each group (control, new1, new2). This structure allows us to proceed with both the ANOVA and the subsequent post-hoc analysis.

```
#create data frame
```

```
data <- data.frame(technique = rep(c("control", "new1", "new2"), each = 10),  
score = c(76, 77, 77, 81, 82, 82, 83, 84, 85, 89,  
81, 82, 83, 83, 83, 84, 87, 90, 92, 93,  
77, 78, 79, 88, 89, 90, 91, 95, 95, 98))
```

```
#view first six rows of data frame
```

```
head(data)
```

```
technique score
```

```
1 control 76
```

```
2 control 77
```

```
3 control 77
```

```
4 control 81
```

```
5 control 82
```

```
6 control 82
```

Inspecting the initial rows confirms that the data frame is correctly structured, with the `technique` column defining the groups and the `score` column containing the quantitative outcome variable. This clean format is essential for the statistical modeling functions we will employ in the upcoming steps.


Step 2: Visualizing Group Differences with Boxplots

Before diving into formal inferential statistics, it is always beneficial to visualize the data. Graphical inspection helps confirm data integrity, assess basic distribution assumptions (like symmetry or outliers), and provide an initial, intuitive understanding of potential differences between the groups.

Boxplots are particularly useful for comparing the central tendency and spread across multiple categories.

The following R code generates side-by-side boxplots, allowing us to visually compare the distribution of exam scores achieved under the control condition, technique 1, and technique 2. The resulting visual output provides a clear contrast in performance among the groups.

```
boxplot(score ~ technique,  
data = data,  
main = "Exam Scores by Studying Technique",  
xlab = "Studying Technique",  
ylab = "Exam Scores",  
col = "steelblue",  
border = "black")
```



The boxplots visually suggest that the mean and median exam scores increase noticeably from the control group to Technique 1, and then further to Technique 2. Technique 2 appears to have the highest scores, while the control group exhibits the lowest. Although this visual evidence is compelling, we must perform the formal statistical test to determine if these observed differences are due to random chance or if they represent a [statistically significantly different](#) effect of the studying techniques.

Step 3: Conducting the One-Way ANOVA

As the required precursor to Dunnett's test, we must first determine if there is an overall effect of the study technique on exam scores using a one-way ANOVA. The ANOVA serves as the omnibus test, checking the global null hypothesis that all group means are equal (i.e., $\mu_{\text{control}} = \mu_{\text{new1}} = \mu_{\text{new2}}$).

If the ANOVA yields a [p-value](#) below the conventional significance level ($\alpha = 0.05$), we reject the global null hypothesis and conclude that differences exist somewhere among the groups. Only upon rejecting this null hypothesis is it statistically appropriate to proceed to the post-hoc analysis.

The following code fits the linear model in R using the `aov()` function and displays the results using `summary()`:

```
#fit the one-way ANOVA model
model <- aov(score ~ technique, data = data)

#view model output
summary(model)

Df Sum Sq Mean Sq F value Pr(>F)
technique 2 211.5 105.73 3.415 0.0476 *
Residuals 27 836.0 30.96
---
Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Reviewing the ANOVA output, the [p-value](#) associated with the `technique` factor is **0.0476**. Since this value is less than the threshold of 0.05, we conclude that the three study techniques do not produce the same average exam score. This significant finding confirms the necessity of performing Dunnett's test to isolate which specific treatment group means differ from the [control group](#).

Step 4: Implementing and Interpreting Dunnett's Test in R

To execute [Dunnett's test](#) in R, we must utilize the `DunnettTest()` function, which is contained within the robust [DescTools](#) library. It is vital to load this library before attempting to call the function.

The syntax for the `DunnettTest()` function is straightforward:

```
DunnettTest(x, g)
```

Where the arguments are defined as:

x: A numeric vector representing the outcome data (in our case, the `score` variable).

g: A vector specifying the categorical group names (the `technique` variable).

The following R code first loads the required library and then applies the function to our dataset, comparing both 'new1' and 'new2' against the implicit control group:

```
#load DescTools library
```

```
library(DescTools)
```

```
#perform Dunnett's Test
```

```
DunnettTest(x=data$score, g=data$technique)
```

Dunnett's test for comparing several treatments with a control :

95% family-wise confidence level

```
$control
```

```
diff lwr.ci upr.ci pval
```

```
new1-control 4.2 -1.6071876 10.00719 0.1787
```

```
new2-control 6.4 0.5928124 12.20719 0.0296 *
```

```
---
```

```
Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1.0
```

The output provides critical statistical details for the two planned comparisons against the control group, maintaining a 95% family-wise confidence level to control for Type I error inflation.

Interpreting the Results of Dunnett's Test

The resulting table compares each treatment group (new1 and new2) against the baseline control group, providing the mean difference (`diff`), the lower and upper bounds of the confidence interval (`lwr.ci` and `upr.ci`), and the adjusted [p-value](#) (`pval`).

Comparison: new1-control

The mean difference in exam scores between New Study Technique 1 and the control group is **4.2** points. The associated p-value is **0.1787**. Since this p-value is greater than 0.05, we conclude that the difference between Technique 1 and the control group is not [statistically significantly different](#).

Comparison: new2-control

The mean difference in exam scores between New Study Technique 2 and the control group is **6.4**

points. Critically, the corresponding p-value is **0.0296**. Because this p-value is less than 0.05 (and is marked with an asterisk), we conclude that Technique 2 produces a statistically significant improvement in mean exam scores compared to the control group.

In summary, the statistical evidence from [Dunnett's test](#) indicates that only studying technique 2 successfully increased student performance beyond what could be attributed to the standard teaching method (the [control group](#)). The teacher should therefore focus on implementing and evaluating Technique 2 further.

Additional Resources

For those interested in exploring related statistical concepts or alternative post-hoc methods, the following resources provide valuable context and detailed instructions:

[An Introduction to the One-Way ANOVA](#)

[How to Conduct a One-Way ANOVA in R](#)

[How to Perform Tukey's Test in R](#)