

Learn How to Perform a One-Way ANOVA in R

Authored by
Mohammed looti

November 9, 2025

RECOMMENDED CITATION

Mohammed looti (2025). *Learn How to Perform a One-Way ANOVA in R*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=14220>

A [One-Way ANOVA](#) (Analysis of Variance) is a fundamental statistical tool used to determine whether there is a [statistically significant difference](#) between the means of three or more independent and unrelated groups. This technique is indispensable in experimental design when comparing the outcomes of several distinct treatments or categories.

This test is referred to as a **one-way** ANOVA because the analysis focuses on how a single categorical factor--the predictor variable--influences the continuous outcome, which is the [response variable](#). The core purpose is to evaluate if the observed variation in the response variable across different groups is greater than the variation observed within the groups themselves.

Note: If the research design required analyzing the impact of two or more predictor variables simultaneously on the response variable, one would typically conduct a [Two-Way ANOVA](#) or a more complex factorial design.

How to Conduct a One-Way ANOVA in R

The following comprehensive example demonstrates the necessary steps for conducting, interpreting, and reporting a One-Way ANOVA using the R statistical environment. We will cover data preparation, preliminary exploration, model fitting, assumption checking, and post hoc analysis.

Background and Data Preparation

Imagine a scenario where we aim to determine if three distinct exercise programs (Program A, Program B, and Program C) yield different average levels of weight loss. In this experiment, the categorical factor we are investigating is the *exercise program* (our predictor variable), and the continuous outcome measure is *weight loss*, recorded in pounds (our response variable).

To address this question, we utilize a One-Way ANOVA to formally test the null hypothesis that the mean weight loss resulting from all three programs is identical. Rejecting this null hypothesis suggests that at least one program results in a significantly different average weight loss compared to the others.

For our study, 90 participants were recruited. We employed a randomized design, assigning 30 people to follow each of the three programs for a period of one month. The following R code simulates the creation of the data frame that we will use for our analysis, ensuring the example is reproducible:

```
#make this example reproducible  
set.seed(0)
```

```
#create data frame
data <- data.frame(program = rep(c("A", "B", "C"), each = 30),
weight_loss = c(runif(30, 0, 3),
runif(30, 0, 5),
runif(30, 1, 7)))

#view first six rows of data frame
head(data)

# program weight_loss
#1 A 2.6900916
#2 A 0.7965260
#3 A 1.1163717
#4 A 1.7185601
#5 A 2.7246234
#6 A 0.6050458
```

The resulting data frame contains two critical columns: the first identifies the specific exercise *program* the participant was assigned to for the month, and the second, *weight_loss*, records the total weight reduction achieved by that individual by the conclusion of the program.

Exploratory Data Analysis (EDA)

Before fitting any formal statistical model, it is essential to conduct exploratory data analysis (EDA). EDA allows us to gain a preliminary understanding of the data structure, identify potential differences between groups, and detect outliers. We begin by calculating descriptive statistics--specifically the mean and standard deviation of weight loss--for each of the three program groups using the powerful **dplyr** package in R.

#load **dplyr** package

```
library(dplyr)
```

```
#find mean and standard deviation of weight loss for each treatment group
```

```
data %>%
group_by(program) %>%
summarise(mean = mean(weight_loss),
sd = sd(weight_loss))
```

```
# A tibble: 3 x 3
```

```
# program mean sd
```

```
#
```


```
#1 A 1.58 0.905  
#2 B 2.56 1.24  
#3 C 4.13 1.57
```

These summary statistics already suggest differences: Program C shows the highest mean weight loss (4.13 lbs), followed by Program B (2.56 lbs), and Program A (1.58 lbs). Furthermore, the standard deviations (SD) indicate that the variability in weight loss is also highest in Program C, which is an important observation when checking model assumptions later.

To visualize the distribution of weight loss across the three programs, we can generate comparative boxplots. Boxplots are excellent tools for comparing group distributions, central tendency (median), spread (interquartile range), and potential skewness or outliers simultaneously.

```
#create boxplots
```

```
boxplot(weight_loss ~ program,  
data = data,  
main = "Weight Loss Distribution by Program",  
xlab = "Program",  
ylab = "Weight Loss",  
col = "steelblue",  
border = "black")
```



The visual evidence from these boxplots strongly supports the summary statistics. We can clearly see that the median weight loss (represented by the line inside the box) for participants in Program C is significantly higher than for the other two groups. Conversely, Program A participants achieved the lowest weight loss. Additionally, the standard deviation, reflected in the overall vertical spread of the boxplot, is noticeably larger for Program C compared to Programs A and B, confirming the quantitative findings regarding variability. The next step is to determine if these visual differences are statistically significant.

Fitting and Interpreting the ANOVA Model

To formally test for significant differences between the group means, we fit the One-Way ANOVA model using R's built-in `aov()` function. This function uses a formula syntax common in linear models across R.

The general syntax to fit a one-way ANOVA model in R is:

```
aov(response variable ~ predictor_variable, data = dataset)
```

For our weight loss example, we designate *weight_loss* as the response variable and *program* as

the predictor variable. After fitting the model object, we use the **summary()** function to generate the standard ANOVA table, which contains the critical test statistics for interpretation:

#fit the one-way ANOVA model

```
model <- aov(weight_loss ~ program, data = data)
```

```
#view the model output
```

```
summary(model)
```

```
# Df Sum Sq Mean Sq F value Pr(>F)
```

```
#program 2 98.93 49.46 30.83 7.55e-11 ***
```

```
#Residuals 87 139.57 1.60
```

```
#---
```

```
#Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Interpreting the ANOVA table is crucial. We focus on the row corresponding to the predictor variable, *program*. The **F value** (30.83) is the ratio of the variance between the groups to the variance within the groups. A large F value suggests that the group means are significantly different. Most importantly, the associated **Pr(>F)** value, which is the [p-value](#), is extremely small (7.55e-11).

Since this p-value is far below the conventional 0.05 significance level, we confidently reject the null hypothesis. This finding confirms that the predictor variable *program* is [statistically significant](#), meaning there is compelling evidence that the mean weight loss varies among the three exercise programs. However, the ANOVA only tells us that a difference exists somewhere; it does not specify which pairs of programs differ from each other.

Validating Model Assumptions and Post Hoc Testing

The reliability of our ANOVA results hinges on whether the underlying [Model Assumptions](#) are met. If these assumptions are severely violated, the p-values and F-statistics derived from the model may be inaccurate. For a One-Way ANOVA, the three critical assumptions are:

Independence: The observations within and between groups must be independent of one another. Since we utilized a randomized design (participants were randomly assigned to programs), this assumption is generally considered met in our scenario.


Normality: The dependent variable (weight loss) must be approximately [normally distributed](#) within each level of the predictor variable (each program group).

Equal Variance (Homogeneity of Variance): The variances of the response variable must be


equal, or at least approximately equal, across all groups.

We can visually inspect the assumptions of normality and equal variance by using the generic `plot()` function on our ANOVA model object, which generates a series of diagnostic plots. We primarily focus on the Residuals vs Fitted plot and the Normal Q-Q plot.

plot(model)



The Q-Q *plot* is used to check the [normality](#) assumption. Ideally, the standardized residuals should align closely with the straight diagonal line. In the plot above, however, the residuals deviate noticeably from the line, particularly at the tails. This visual evidence suggests a potential, albeit mild, violation of the normality assumption.



The *Residuals vs Fitted* plot helps us assess the assumption of [Equal Variance](#) (Homoscedasticity). For the assumption to be met, the residuals should be equally scattered above and below the horizontal line across all fitted values. In our case, the spread of the residuals is notably greater for higher fitted values, indicating that the variance is not constant across all programs. This suggests that the assumption of [Equal Variance](#) may be violated.

To formally test the homogeneity of variances, we perform [Levene's Test](#), which requires the **car** package:

```
#load car package
```

```
library(car)
```

```
#conduct Levene's Test for equality of variances
```

```
leveneTest(weight_loss ~ program, data = data)
```

```
#Levene's Test for Homogeneity of Variance (center = median)
```

```
# Df F value Pr(>F)
```

```
#group 2 4.1716 0.01862 *
```

```
# 87
```

```
#---
#Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

The [p-value](#) from Levene's Test is **0.01862**. Using the standard 0.05 significance level, we would reject the null hypothesis of equal variances, confirming the violation suggested by the residuals plot. While significant violations of assumptions might necessitate data transformation or the use of non-parametric tests, for the purpose of this demonstration, we will proceed with the post hoc analysis, acknowledging the potential limitation.

Since the overall ANOVA indicated a [statistically significant difference](#) among the groups, the next logical step is to perform a [Post Hoc Test](#). Post hoc tests are necessary to pinpoint exactly which specific pairs of group means are different. We will use the function **TukeyHSD()** to conduct [Tukey's Honestly Significant Difference \(HSD\) Test](#) for multiple comparisons, as it controls the family-wise error rate.


```
#perform Tukey's Test for multiple comparisons
TukeyHSD(model, conf.level=.95)
```

```
# Tukey multiple comparisons of means
# 95% family-wise confidence level
#
#Fit: aov(formula = weight_loss ~ program, data = data)
#
# $program
# diff lwr upr p adj
#B-A 0.9777414 0.1979466 1.757536 0.0100545
#C-A 2.5454024 1.7656076 3.325197 0.0000000
#C-B 1.5676610 0.7878662 2.347456 0.0000199
```

The output of [Tukey's Test](#) provides pairwise comparisons (B-A, C-A, C-B), the mean difference (diff), the 95% confidence interval bounds (lwr and upr), and the adjusted p-value (p adj). Since all adjusted p-values are below the 0.05 threshold (0.010, 0.000, and 0.00002), we conclude that there is a significant difference in mean weight loss between every single pair of programs.

We can further visualize these differences and their confidence intervals using the ``plot(TukeyHSD())`` function:

```
#create confidence interval for each comparison
plot(TukeyHSD(model, conf.level=.95), las = 2)
```



This plot provides a graphical confirmation of the hypothesis test results. A critical visual check is whether the horizontal confidence interval for any comparison crosses the zero vertical line. Since none of the 95% confidence intervals for the mean differences between the programs include the value zero, this confirms that all three pairwise comparisons demonstrate a [statistically significant difference](#) in mean weight loss.

Reporting the Results of the One-Way ANOVA

The final step in any statistical analysis is to clearly and concisely report the findings, typically adhering to established guidelines (such as APA format in many fields), which summarizes the overall ANOVA result and the subsequent post hoc comparisons.

A One-Way ANOVA was conducted to examine the effects of the exercise program (A, B, or C) on weight loss, measured in pounds. The results indicated a [statistically significant difference](#) between the effects of the three programs on weight loss ($F(2, 87) = 30.83, p = 7.55e-11$). The effect size was large, suggesting that the program selection accounts for a substantial proportion of the variance in weight loss.

Following the significant omnibus F-test, [Tukey's HSD post hoc tests](#) were carried out to determine specific group differences, controlling for the family-wise error rate. The results confirmed that every program yielded a significantly different mean weight loss outcome from the others.

Specifically, the mean weight loss for participants in Program C (M=4.13, SD=1.57) was significantly higher than the mean weight loss for participants in Program B (M=2.56, SD=1.24; $p < 0.0001$). Furthermore, the mean weight loss for participants in Program C was also significantly higher than for participants in Program A (M=1.58, SD=0.91; $p < 0.0001$).

In addition, the mean weight loss observed for Program B participants was significantly higher than the mean weight loss for Program A participants ($p = 0.010$). These findings collectively suggest that Program C is the most effective intervention for weight loss among the three tested options.

Additional Resources

For those seeking to deepen their understanding of hypothesis testing and ANOVA methodologies, the following tutorials provide additional information and related statistical concepts: