

How to Perform a One-Way ANOVA in R: A Step-by-Step Guide

Authored by
stats writer

March 3, 2026

RECOMMENDED CITATION

stats writer (2026). *How to Perform a One-Way ANOVA in R: A Step-by-Step Guide*. PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=133684>

Understanding the Fundamentals of One-Way ANOVA in R

The **One-Way Analysis of Variance**, commonly referred to as **One-Way ANOVA**, is a foundational statistical technique employed to determine whether significant differences exist between the means of three or more independent groups. This method is particularly effective when researchers need to assess the impact of a single categorical independent variable on a continuous dependent variable. By analyzing the variance within and between these groups, the **ANOVA** test evaluates the null hypothesis, which posits that all group means are equal, against an alternative hypothesis suggesting that at least one group mean differs from the others.

The "one-way" designation signifies that the analysis involves only one factor or predictor variable. If a researcher were interested in examining the simultaneous effects of two different categorical variables on a single response, they would instead utilize a Two-Way ANOVA. In the R programming language, conducting a **One-Way ANOVA** is streamlined through the use of built-in functions designed for linear modeling and variance decomposition, making it a preferred tool for data scientists and statisticians alike.

Mastering the implementation of **ANOVA** in **R** requires an understanding of both the mathematical underpinnings of the test and the specific syntax required to execute it. This involves preparing a structured **data frame**, selecting the appropriate model function--typically **aov()**--and interpreting the resulting **p-values** and **F-statistics**. Beyond the initial test, a comprehensive analysis also includes post-hoc testing and assumption validation to ensure that the conclusions drawn from the data are both accurate and reliable.

Experimental Design and Data Preparation

To illustrate the practical application of a **One-Way ANOVA**, consider a hypothetical study designed to measure the efficacy of three distinct exercise programs on weight loss. In this scenario, the predictor variable is the **exercise program** (with levels A, B, and C), and the response variable is the total weight loss recorded in pounds. This experimental setup allows us to explore how different interventions influence a physical health outcome across a diverse cohort of participants.

In this example, we recruit 90 individuals and randomly assign 30 participants to each of the three programs. This **randomization** is critical as it helps mitigate the influence of confounding variables, ensuring that any observed differences in weight loss can be more confidently attributed to the exercise programs themselves. To maintain **reproducibility** in our **R** environment, we use the **set.seed()** function, which ensures that the random numbers generated for our dataset remain consistent across different sessions.

The following **R** code block demonstrates how to construct the necessary **data frame** for this

analysis. We utilize the **rep()** function to create group labels and the **runif()** function to simulate weight loss data within specific ranges for each group, reflecting the expected variations in program performance:

```
#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
```

Proper data structuring is essential for the **aov()** function to interpret the relationship between variables correctly. In the code above, the **program** column acts as the categorical factor, while the **weight_loss** column provides the numerical measurements. Ensuring that the categorical variable is stored as a **factor** in **R** is a common best practice, as it prevents the software from treating group labels as continuous numerical data during the modeling process.

Exploratory Data Analysis and Visualization

Prior to fitting a formal **statistical model**, it is imperative to conduct an **Exploratory Data Analysis (EDA)**. This phase allows the researcher to gain an intuitive understanding of the data's distribution and central tendency. By calculating the **mean** and **standard deviation** for each group, we can identify preliminary patterns and potential outliers. The **dplyr** package is an excellent tool for this purpose, providing a readable syntax for grouping and summarizing datasets.

```
#load dplyr packagelibrary(dplyr)
```

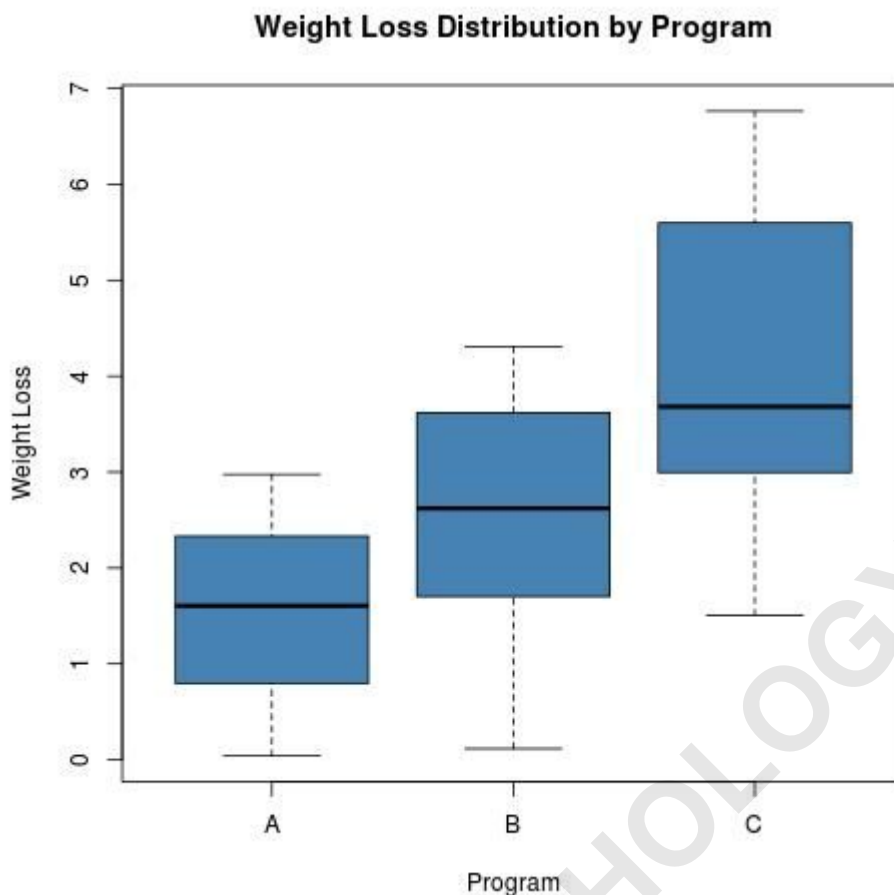
```
#find mean and standard deviation of weight loss for each treatment group
```

```
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
```

Numerical summaries are often best complemented by **data visualization**. A **boxplot** is particularly useful in the context of **ANOVA**, as it highlights the **median**, **interquartile range**, and overall spread of the data for each category. Observing the overlap--or lack thereof--between boxes provides an immediate visual cue regarding the likelihood of significant differences between the group means. In our example, a **boxplot** will help us visualize the distribution of weight loss across Programs A, B, and C.

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



The resulting visual output clearly indicates that participants in Program C generally experienced higher weight loss than those in Programs A or B. However, we also observe that Program C exhibits a wider **variance** (represented by the length of the box and whiskers). While these visual trends are informative, they do not provide **statistical significance**; therefore, we must proceed with the **ANOVA** model to determine if these differences are mathematically meaningful or simply the result of random variation.

Executing the One-Way ANOVA Model in R

To perform the actual **One-Way ANOVA**, R utilizes the **aov()** function, which is designed to fit a linear model to the data. The syntax for this function follows the standard R formula notation: **response ~ predictor**. This translates to "weight loss is a function of the exercise program." By specifying the **data frame** within the function call, R can efficiently map the variables to their respective values.

After fitting the model, the **summary()** function is used to extract the **ANOVA table**. This table is the centerpiece of the analysis, containing the **Sum of Squares**, **Mean Squares**, and the **F-value**. The most critical component for many researchers is the **Pr(>F)** value, which represents the **p-**

value. A **p-value** below a predetermined threshold (typically 0.05) indicates that we should reject the **null hypothesis** in favor of the **alternative hypothesis**.

#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
```

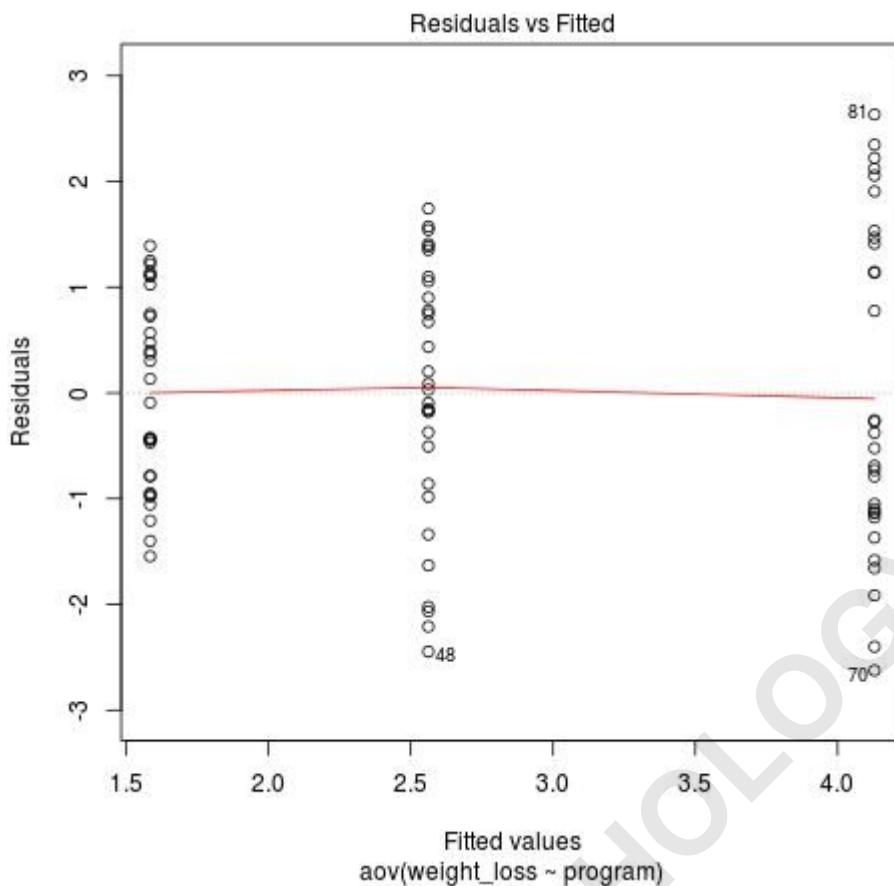
In our current analysis, the **p-value** is extremely small ($7.55e-11$), which is far below the standard 0.05 significance level. This leads us to conclude that the type of exercise program has a statistically significant effect on weight loss. The **F-statistic** of 30.83 further supports this conclusion, as a high **F-value** suggests that the variation between the group means is significantly larger than the variation within the groups themselves. However, the **ANOVA** result does not specify which specific groups are different; it only tells us that at least one pair is distinct.

Validating Statistical Assumptions

For the results of a **One-Way ANOVA** to be considered valid, several **statistical assumptions** must be satisfied. The first is **independence**, meaning that the data points in one group should not be influenced by or related to those in another group. In our experimental design, this is achieved through random assignment. The second major assumption is **normality**, which requires that the **residuals** of the model follow a **normal distribution**. If the data is heavily skewed, the reliability of the **p-value** may be compromised.

The third assumption is **homogeneity of variance**, also known as **homoscedasticity**. This implies that the different groups should have approximately equal variances. When variances are drastically different, the **ANOVA** test may become either too conservative or too liberal, leading to **Type I or Type II errors**. We can evaluate these assumptions visually using **diagnostic plots** generated by the **plot()** function in R, which provides a **Residuals vs Fitted** plot and a **Normal Q-Q** plot.

```
plot(model)
```

In our diagnostic plot, the residuals for higher fitted values appear more spread out than those for lower values, suggesting that the variance is not constant across groups. To move beyond visual estimation, we can perform a formal Levene's Test using the **car package**. This test evaluates the null hypothesis that the variances are equal across all groups.

```
#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 **Levene's Test** yields a **p-value** of 0.01862. If we adhere to a strict 0.05 significance level, we would reject the assumption of equal variances. However, if our study requires a more stringent 0.01 level, the assumption might still be considered acceptable. When variances are significantly unequal, statisticians often turn to **Welch's ANOVA**, which does not require the assumption of **homoscedasticity**, ensuring the integrity of the analysis even with irregular data spreads.

Post-Hoc Analysis with Tukey's HSD

Once a **One-Way ANOVA** confirms that significant differences exist among group means, the next logical step is to identify exactly which groups differ from one another. This is where **post-hoc tests** become essential. A simple series of **t-tests** would not be appropriate here because of the **multiple comparisons problem**, which increases the likelihood of committing a **Type I error** (false positive). To control for this, we use **Tukey's Honestly Significant Difference (HSD)** test.

The **TukeyHSD()** function in **R** calculates the pairwise differences between all possible combinations of groups and adjusts the **p-values** to maintain the **family-wise error rate**. This provides a more rigorous and accurate assessment of which exercise programs are truly different in their impact on weight loss. The output includes the mean difference, the lower and upper bounds of the **95% confidence interval**, and the adjusted **p-value** for each comparison.

#perform Tukey's Test for multiple comparisons

```
TukeyHSD(model, conf.level=.95)
```

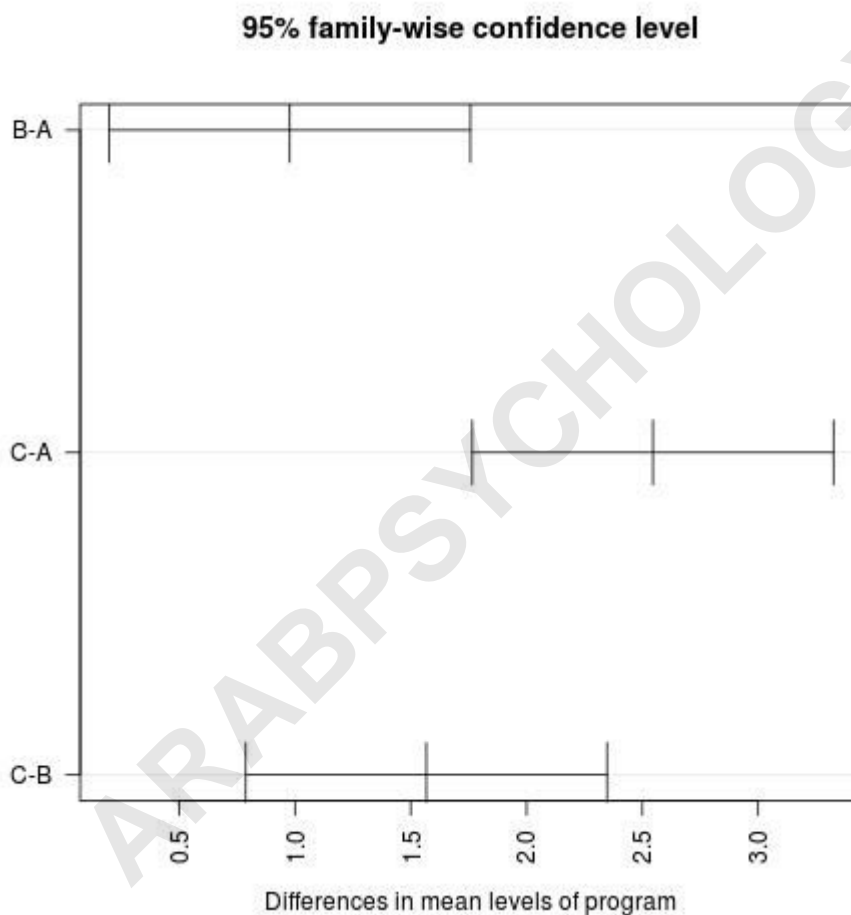
```
# Tukey multiple comparisons of means
# 95% family-wise confidence level
#
#Fit: aov(formula = weight_loss ~ program, data = data)
#
# $program
# diff lwr upr p adj
#B-A 0.9777414 0.1979466 1.757536 0.0100545
#C-A 2.5454024 1.7656076 3.325197 0.0000000
#C-B 1.5676610 0.7878662 2.347456 0.0000199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

Based on the **Tukey HSD** results, every pairwise comparison is statistically significant. For instance, the difference between Program B and Program A is approximately 0.98 pounds ($p = 0.01$), while the difference between Program C and Program A is a much larger 2.55 pounds ($p < 0.001$). These results allow us to rank the programs effectively, concluding that Program C is the most effective for weight loss, followed by Program B, and then Program A.

Visualizing Pairwise Differences and Confidence Intervals

To make the results of the **post-hoc analysis** more accessible, **R** allows for the direct plotting of the **Tukey HSD** output. This visualization displays the **confidence intervals** for the differences between each pair of means. If a confidence interval for a pair does not cross the vertical line at zero, it signifies that the difference between those two groups is statistically significant at the chosen confidence level (usually 95%).

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



The **Tukey HSD plot** serves as a powerful communication tool for summarizing complex **statistical data**. In our example, all three horizontal lines representing the comparisons (B-A, C-A, and C-B) are positioned entirely to the right of the zero line. This visual evidence confirms that all programs are significantly different from each other, reinforcing the numerical findings from the **ANOVA** table and the **TukeyHSD()** output.

Furthermore, these plots help researchers identify the **magnitude** of the differences. For example, the gap between Program C and Program A is much larger than the gap between Program B and Program A, which is immediately apparent by the position of the intervals along the x-axis. Such visual insights are invaluable when presenting findings to stakeholders who may not be deeply familiar with **p-values** or **F-statistics** but can easily interpret interval charts.

Synthesizing and Reporting the Research Findings

The final stage of any statistical analysis is the clear and concise reporting of the results. In academic and professional settings, it is standard practice to report the **F-statistic**, the **degrees of freedom**, and the **p-value**. Additionally, the results of **post-hoc tests** should be included to provide a complete picture of the group relationships. This ensures that the reader understands not only that a difference exists, but also the nature and direction of that difference.

A formal summary for our exercise program study might be written as follows: A **One-Way ANOVA** was conducted to evaluate the effect of three different exercise programs on weight loss. The analysis revealed a statistically significant difference between the programs ($F(2, 87) = 30.83, p < 0.001$). Following this, **Tukey's HSD post-hoc tests** were performed to identify specific group differences. The mean weight loss for Program C was significantly higher than both Program B ($p < 0.001$) and Program A ($p < 0.001$). Furthermore, Program B was found to be significantly more effective than Program A ($p = 0.01$).

By following this structured approach--from data preparation and **Exploratory Data Analysis** to model fitting, assumption checking, and post-hoc testing--researchers can utilize **R** to perform robust and reproducible **One-Way ANOVA**. This workflow not only ensures the mathematical validity of the conclusions but also provides the necessary tools to communicate those findings effectively to a broader audience. For those looking to expand their knowledge, further study into **Two-Way ANOVA**, **ANCOVA**, and **linear mixed-effects models** in **R** is highly recommended.