

How to Perform an ANCOVA in R: A Step-by-Step Guide

Authored by
stats writer

March 3, 2026

RECOMMENDED CITATION

stats writer (2026). *How to Perform an ANCOVA in R: A Step-by-Step Guide*.

PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=133634>

ANCOVA (Analysis of Covariance) is a statistical method used to compare the means of two or more groups while controlling for the influence of one or more continuous variables, known as covariates. In R, ANCOVA can be conducted by using the "aov" function, which stands for Analysis of Variance, and the "lm" function, which stands for Linear Model. First, the data needs to be imported into R and the variables of interest need to be identified. Then, the "lm" function is used to create a linear model by specifying the dependent variable, independent variable, and covariates. The "aov" function is then used to fit the model and perform the ANCOVA analysis. The results can be interpreted and visualized using various tools and techniques in R. Overall, ANCOVA in R allows for a systematic and efficient way to analyze the relationship between multiple groups while considering the influence of covariates.

Conduct an ANCOVA in R

This tutorial provides an example of how to perform an ANCOVA in R.

Example: ANCOVA in R

We will conduct an ANCOVA to test whether or not studying technique has an impact on exam scores by using the following variables:

**Studying technique: The independent variable we are interested in analyzing
Student's current grade: The covariate that we want to take into account
Exam score: The response variables we are interested in analyzing**

The following dataset contains information for 90 students that were randomly split into three groups of 30.

The dataset shows the studying technique each student used (*A*, *B*, or *C*), their current grade in the class when they started using the studying technique, and their exam score they received after using the studying technique for one month to prepare for the exam:

```
#make this example reproducible  
set.seed(10)
```

```
#create dataset
```

```
data <- data.frame(technique = rep(c("A", "B", "C"),  
each = 30),  
current_grade = runif(90, 65, 95),  
exam = c(runif(30, 80, 95), runif(30, 70, 95), runif(30, 70,  
90)))
```

```
#view first six lines of dataset
```

```
head(data)
```

```
# technique current_grade exam
```

```
#1 A 80.22435 87.32759
```

```
#2 A 74.20306 90.67114
```

```
#3 A 77.80723 88.87902
```

```
#4 A 85.79306 87.75735
```

```
#5 A 67.55408 85.72442
```

#6 A 71.76310 92.52167

Step 1: Explore the Data

Before we fit the ANCOVA model, we should first explore the data to gain a better understanding of it and verify that there aren't any extreme outliers that could skew the results.

First, we can view a summary of each variable in the dataset:

```
summary(data)
```

```
# technique current_grade exam  
# A:30 Min. :65.43 Min. :71.17  
# B:30 1st Qu.:71.79 1st Qu.:77.27  
# C:30 Median :77.84 Median :84.69  
# Mean :78.15 Mean :83.38  
# 3rd Qu.:83.65 3rd Qu.:89.22  
# Max. :93.84 Max. :94.76
```

We can see that each value for studying technique (A, B, and C) shows up 30 times in the data.

We can also see how the current student scores were

distributed at the beginning of the study. The minimum score in the class was 65.43, the max was 93.84, and the mean was 78.15.

Likewise, we can see that the minimum score received on the exam was 71.17, the max was 94.76, and the mean was 83.38.

Next, we can use the `dplyr` package to easily find the mean and the standard deviation of both the current grades and the exam scores for each studying technique:

```
#load dplyr  
library(dplyr)  
  
data %>%  
group_by(technique) %>%  
summarise(mean_grade = mean(current_grade),  
sd_grade = sd(current_grade),  
mean_exam = mean(exam),  
sd_exam = sd(exam))  
  
# A tibble: 3 x 5  
# technique mean_grade sd_grade mean_exam
```

```
sd_exam
```

```
#1 A 79.0 7.00 88.5 3.88
```

```
#2 B 78.5 8.33 81.8 7.62
```

```
#3 C 76.9 8.24 79.9 5.71
```

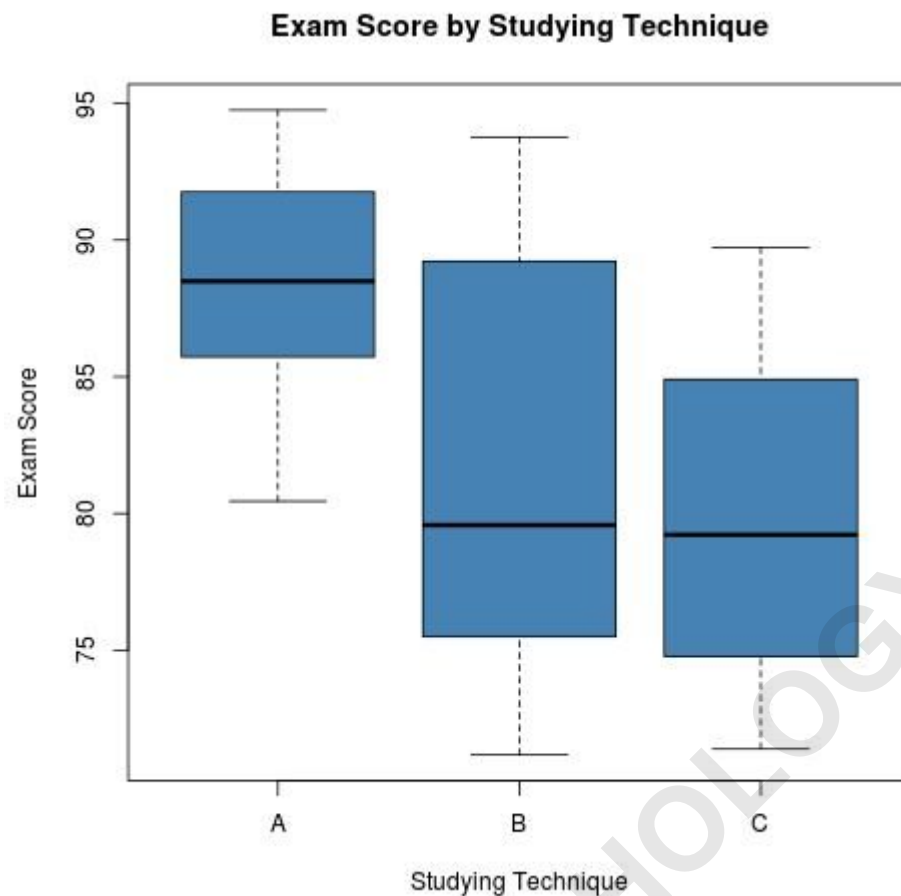
We can see that the mean and the standard deviations of the current grade for the students using each studying technique is roughly similar.

We can also visualize the distribution of exam scores based on studying technique by using boxplots:

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

ARABPSYCHOLOGY.COM


ARABPSYCHOLOGY.COM



Similarly, we can also use boxplots to visualize the distribution of *current grades* based on studying technique:

```
boxplot(current_grade ~ technique,  
data = data,  
main = "Current Grade by Studying Technique",  
xlab = "Studying Technique",  
ylab = "Current Grade",  
col = "steelblue",  
border = "black")
```

)



Step 2: Check the Model Assumptions

Once we've done some basic data exploration and are familiar with the data, we need to verify that the following assumptions for ANCOVA are met:

The covariate and the treatment are independent - we need to verify that the covariate (*current grade*) and the treatment (*studying technique*) are independent of each

other, since adding a covariate term into the model only makes sense if the covariate and the treatment act independently on the response variable (*exam*). Homogeneity of variance - we need to verify that the variances among the groups is equal

To verify that the covariate and the treatment are independent, we can run an ANOVA using *current grade* as the response variable and *studying technique* as the predictor variable:

```
#fit anova model
anova_model <- aov(current_grade ~ technique, data =
data)
#view summary of anova model
summary(anova_model)

# Df Sum Sq Mean Sq F value Pr(>F)
#technique 2 74 37.21 0.599 0.552
#Residuals 87 5406 62.14
```

The p-value is greater than 0.05, so the covariate (*current grade*) and the treatment (*studying technique*) seem to be independent.

Next, to verify that there is homogeneity of variance among the groups, we can conduct Levene's Test:

```
#load car library to conduct Levene's Test
```

```
library(car)
```

```
#conduct Levene's Test
```

```
leveneTest(exam~technique, data = data)
```

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

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

```
#group 2 9.4324 0.0001961 ***
```

```
# 87
```

The p-value from the test is equal to .0001961, which indicates that the variances among the groups are not equal. Although we could attempt a transformation on the data to correct this problem, we won't worry too much about the differences in variance for the time being.

Step 3: Fit the ANCOVA Model

Next, we'll fit the ANCOVA model using *exam score* as the response variable, *studying technique* as the

predictor (or "treatment") variable, and *current grade* as the covariate.

We'll use the `Anova()` function in the `car` package to do so, just so we can specify that we'd like to use type III sum of squares for the model, since type I sum of squares is dependent upon the order that the predictors are entered into the model:

```
#load car library  
library(car)
```

```
#fit ANCOVA model
```

```
ancova_model <- aov(exam ~ technique +  
current_grade, data = data)
```

```
#view summary of model
```

```
Anova(ancova_model, type="III")
```

```
#Response: exam
```

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

```
 #(Intercept) 7161.2 1 201.4621 < 2.2e-16 ***
```

```
 #technique 1242.9 2 17.4830 4.255e-07 ***
```

```
 #current_grade 12.3 1 0.3467 0.5576
```

```
 #Residuals 3057.0 86
```

We can see that the p-value for *technique* is extremely small, which indicates that studying technique has a statistically significant effect on exam scores, even after controlling for the current grade.

Step 4: Post Hoc Tests

Although the ANCOVA results told us that *studying technique* had a statistically significant effect on exam scores, we need to run post hoc tests to actually find out which studying techniques differ from each other.

To do so, we can use the `glht()` function within the `multcomp` package in R to perform Tukey's Test for multiple comparisons:

```
#load the multcomp library
```

```
library(multcomp)
```

```
#fit the ANCOVA model
```

```
ancova_model <- aov(exam ~ technique +  
current_grade, data = data)
```

```
#define the post hoc comparisons to make
```

```
postHocs <- glht(ancova_model, linfct = mcp(technique  
= "Tukey"))
```

```
#view a summary of the post hoc comparisons  
summary(postHocs)
```

```
#Multiple Comparisons of Means: Tukey Contrasts
```

```
#
```

```
#Fit: aov(formula = exam ~ technique + current_grade,  
data = data)
```

```
#
```

```
#Linear Hypotheses:
```

```
# Estimate Std. Error t value Pr(>|t|)
```

```
#B - A == 0 -6.711 1.540 -4.358 0.000109 ***
```

```
#C - A == 0 -8.736 1.549 -5.640 < 1e-04 ***
```

```
#C - B == 0 -2.025 1.545 -1.311 0.393089
```

```
#view the confidence intervals associated with the  
multiple comparisons  
confint(postHocs)
```

```
# Simultaneous Confidence Intervals
```

```
#
```

```
#Multiple Comparisons of Means: Tukey Contrasts
```

```
#
```

```
#Fit: aov(formula = exam ~ technique + current_grade,  
data = data)
```

```
#
```

```
#Quantile = 2.3845
#95% family-wise confidence level
#
#Linear Hypotheses:
# Estimate lwr upr
#B - A == 0 -6.7112 -10.3832 -3.0392
#C - A == 0 -8.7364 -12.4302 -5.0426
#C - B == 0 -2.0252 -5.7091 1.6588
```

From the output, we can see that there is a statistically significant difference (at $\alpha = .05$) in exam scores between studying technique *A* and studying technique *B* (p-value: .000109) as well as between technique *A* and technique *C* (p-value: $<1e-04$).

We can also see that there is *not* a statistically significant difference (at $\alpha = .05$) between techniques *B* and *C*. The confidence intervals between the techniques confirm these conclusions as well.

Thus, we can conclude that using studying technique *A* leads to a statistically significantly greater exam score for students compared to techniques *B* and *C*, even after controlling for the student's current grade in the

class.

ARABPSYCHOLOGY.COM