

A Step-by-Step Guide to Analysis of Covariance (ANCOVA) with Python

Authored by
Mohammed loot

November 8, 2025

RECOMMENDED CITATION

Mohammed loot (2025). *A Step-by-Step Guide to Analysis of Covariance (ANCOVA) with Python*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=12686>

The [Analysis of Covariance \(ANCOVA\)](#) stands as a sophisticated statistical technique essential for researchers aiming to isolate the true effect of a categorical factor on a dependent variable. It is specifically designed to determine if statistically significant differences exist between the means of multiple independent groups, all while systematically accounting for the influence of one or more continuous variables. These control variables are formally known as [covariates](#). Fundamentally, ANCOVA expertly merges components from traditional [Analysis of Variance \(ANOVA\)](#)--which handles categorical grouping--with the precision of linear regression, allowing analysts to remove extraneous variance explained by continuous external factors.

This comprehensive, expert-level guide is dedicated to detailing the step-by-step process of executing a robust ANCOVA model within the modern Python data science ecosystem. We will focus on leveraging powerful libraries such as **pandas** for data manipulation and the specialized **pingouin** package for seamless statistical modeling. Our exploration moves beyond mere execution; we aim to solidify your understanding of the underlying statistical logic and the proper methodology for interpreting the resulting output, ensuring you can confidently and effectively deploy this technique in complex data analysis and experimental research projects.

The Theoretical Foundation: Why ANCOVA is Necessary

ANCOVA plays a critical role in enhancing the rigor of experimental design, particularly in real-world scenarios where true random assignment of subjects or perfect control over initial conditions is impossible. By integrating a [covariate](#) into the model, ANCOVA achieves a statistical adjustment: it modifies the estimated means of the dependent variable as if every subject had scored identically on the covariate. This powerful statistical equalization significantly increases the precision and reliability of the overall analysis.

The primary and most compelling advantage of utilizing [ANCOVA](#) is its capacity to substantially reduce the error variance, often referred to as the within-group variance or residual variance. When a selected continuous variable is demonstrably correlated with the response variable, including it as a covariate allows the model to partition out the variance that this variable accounts for. This process decreases the residual sum of squares, resulting in a cleaner and more sensitive statistical test. A reduction in unexplained noise means the model is far more likely to detect genuine group differences if they exist, thereby maximizing the **statistical power** of the experimental findings.

To properly construct and understand an ANCOVA model, one must distinguish its core components: first, the categorical independent variable, known as the factor, which defines the groups being compared; second, the continuous dependent variable, or the response variable, which is the outcome being measured; and finally, the one or more continuous independent variables, the covariates, which serve as controls. The central objective of the analysis is to ascertain whether the categorical factor still exerts a significant influence on the response variable

after the linear effect of the [covariate](#) has been statistically removed or held constant across the groups.

ANCOVA Distinguished from ANOVA and Regression

Contextualizing ANCOVA through comparison with its related techniques--ANOVA and linear regression--is vital for understanding its specialized utility. Standard [ANOVA](#) is optimally suited for comparing group means based exclusively on categorical factors. However, if there are continuous variables that significantly predict the outcome and differ systematically across groups, ignoring them can confound the results, leading to misleading conclusions due to inflated noise or unexplained variance within groups.

Conversely, while standard multiple linear regression can technically accommodate both categorical (through dummy coding) and continuous predictors, the structure of ANCOVA is specifically engineered to test the main effect of the categorical factor while rigorously controlling for the continuous variable. Crucially, ANCOVA yields adjusted means for each group. These adjusted means represent the expected outcome for each group if all subjects were statistically equivalent on the covariate, thus making the comparison between groups significantly fairer and more direct by neutralizing pre-existing differences.

This subtle but critical distinction makes ANCOVA indispensable in fields like psychology, clinical trials, and educational research, where initial conditions are rarely perfectly balanced. For instance, if a study is evaluating a new therapeutic intervention, and one treatment group inadvertently possesses higher baseline disease severity scores (the covariate), ANCOVA allows researchers to statistically equalize these groups. The analysis then adjusts the final outcome scores, ensuring that any observed effect of the intervention is a true reflection of the treatment, not merely an artifact of those initial, pre-existing differences in baseline metrics.

Practical Application: A Case Study in Educational Research

To fully grasp the practical necessity of ANCOVA in Python, let us examine a typical educational research scenario. Imagine a school administrator investigating the efficacy of three distinct studying techniques--labeled A, B, and C--on students' final exam scores. The administration acknowledges that a student's prior academic performance, specifically their current grade in the course, is a powerful predictor of their final exam success. Therefore, to ensure that the measured impact of the studying technique is purely the result of the intervention itself, the administrator must statistically control for the influence of the initial current grade.

This requirement mandates the setup of an ANCOVA model. The core analytical question guiding this study is: Does the assigned studying technique significantly affect the final exam score, specifically *after* the linear influence of the student's pre-existing current grade has been

statistically accounted for and removed from the variance?

The variables within this specific educational study are rigorously defined according to the ANCOVA framework:

Factor variable: This is the categorical independent variable--the three different **studying techniques** (A, B, C). This is the primary variable of interest whose effect we wish to measure.

Covariate: This is the continuous variable we must control for--the student's **current grade** in the class. This variable is anticipated to correlate strongly with the response variable.

Response variable: This is the continuous dependent variable--the final **exam score**. This is the outcome measure used to assess the effectiveness of the techniques.

Step 1: Data Preparation and Environment Setup

The foundational step for any statistical analysis in Python involves setting up the computational environment and preparing the data into a structured format. We rely on the **NumPy** library for efficient numerical operations and the powerful [pandas](#) library to construct a structured DataFrame. The DataFrame organizes the experimental variables into columns, making the data easily readable and accessible for the specialized statistical modeling functions provided by the **pingouin** library in subsequent steps.

The following code block demonstrates the standard procedure for importing the necessary libraries and constructing the dataset. Each observation, represented by a row, contains the assigned technique, the continuous current grade (the covariate), and the resulting exam score (the response). We utilize a compact, balanced dataset consisting of 15 observations, where five students were assigned to each of the three studying techniques.

```
import numpy as np
import pandas as pd
```

```
#create data
df = pd.DataFrame({'technique': np.repeat(, 5),
'current_grade': ,
'exam_score': })
#view data
df
```

```
technique current_grade exam_score
0 A 67 77
1 A 88 89
2 A 75 72
```

3 A 77 74
 4 A 85 69
 5 B 92 78
 6 B 69 88
 7 B 77 93
 8 B 74 94
 9 B 88 90
 10 C 96 85
 11 C 91 81
 12 C 88 83
 13 C 82 88
 14 C 80 79

Step 2 & 3: Model Execution and Interpretation of Results

To execute complex statistical tests like ANCOVA within Python, the [Pingouin](#) statistical package is highly recommended for its clear syntax and comprehensive statistical reporting. Assuming Pingouin is installed, the core function for this analysis is `ancova()`. This function requires the analyst to specify the roles of the data variables: `data` (the pandas DataFrame), `dv` (the dependent variable, `exam_score`), `covar` (the continuous variable to control for, `current_grade`), and `between` (the categorical factor defining the groups, `technique`).

Executing the function instructs Python to statistically perform the Analysis of Covariance. The model first adjusts the variance in the exam scores based on the linear relationship with the current grade, and subsequently tests for mean differences between the three study techniques based on these adjusted scores. The output table, generated below, provides the standardized summary required for drawing scientific conclusions.

pip install pingouin

```
from pingouin import ancova
```

```
#perform ANCOVA
```

```
ancova(data=df, dv='exam_score', covar='current_grade', between='technique')
```

```
Source SS DF F p-unc np2
0 technique 390.575130 2 4.80997 0.03155 0.46653
1 current_grade 4.193886 1 0.10329 0.75393 0.00930
2 Residual 446.606114 11 NaN NaN NaN
```

The resulting table is crucial for assessing the independent effects of both the factor and the

covariate. Each row identifies a "Source" of variance accounted for in the model. The key metrics for interpretation include the **F** (F-statistic) and the **p-unc** (the uncorrected [p-value](#)). The [p-value](#) is the metric used to rigorously determine whether the observed results are statistically significant relative to a pre-defined alpha level, typically set at 0.05.

Interpreting the Factor (Studying Technique)

We first examine the row labeled "technique," which reflects the main effect of the three study methods after the variance attributed to the students' current grades has been statistically controlled. In this output, the p-unc value is **0.03155**. Since this value is less than the standard significance threshold ($\alpha = 0.05$), we conclude that the effect of the studying technique on the final exam scores is statistically significant.

This finding of significance compels us to reject the [null hypothesis](#), which stated that all three studying techniques yield the same average exam score. The result confirms that, even after neutralizing initial differences in student academic ability, at least one of the studying techniques performs significantly differently from the others. This robust conclusion validates the intervention's efficacy; however, further post-hoc testing would be required to pinpoint precisely which specific pairs of techniques differ from each other.

Interpreting the Covariate (Current Grade)

Finally, the row labeled "current_grade" evaluates the linear relationship between the covariate and the dependent variable, adjusted for the categorical factor. In this case, the p-unc for the current grade is **0.75393**. Given that this value is considerably greater than 0.05, we do not find a significant effect of the current grade on the final exam score *within the context of this specific ANCOVA model*. While counter-intuitive in an educational setting where current grades usually predict outcomes, this result suggests that after partitioning the variance associated with the grouping factor (technique), the remaining variance explained by the current grade was minimal.

In conclusion, the ANCOVA successfully isolated the unique influence of the study technique. The analysis demonstrated that the technique's effectiveness was significant, regardless of the students' initial academic standing as measured by their current grades, thereby providing a clear and scientifically robust conclusion regarding the intervention's success.