

How to Perform a Nested ANOVA: A Step-by-Step Guide

Authored by
stats writer

December 6, 2025

RECOMMENDED CITATION

stats writer (2025). *How to Perform a Nested ANOVA: A Step-by-Step Guide*.
PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=106269>

The Analysis of Variance (ANOVA) is a fundamental statistical technique used across various scientific disciplines to determine whether there are statistically significant differences between the means of three or more independent groups. While standard ANOVA models, such as the one-way or two-way design, are frequently employed, complex experimental structures often necessitate specialized variations. The **Nested ANOVA**, also referred to as a **Hierarchical ANOVA**, addresses scenarios where experimental factors are arranged in a hierarchical or subset relationship rather than a fully crossed design.

This powerful analytical tool is essential when dealing with multi-stage sampling or experimental designs where the levels of one factor are uniquely contained within the levels of a higher-level factor. Unlike traditional designs that assume all factors interact fully, the Nested ANOVA explicitly models the variance associated with factors that operate within specific contexts or levels of another treatment. This structure ensures accurate partitioning of the variance and provides robust conclusions regarding the effects of both the main and subordinate factors on the response variable.

Understanding the Core Concept: Analysis of Variance (ANOVA)

Before delving into the specifics of nesting, it is helpful to recall the primary objective of any ANOVA: partitioning the total observed variability in the data into components attributable to systematic effects (the factors) and components attributable to random error (residual variation). The F-test statistic derived from this partitioning allows researchers to test null hypotheses concerning the equality of population means across different experimental groups. The precision of these conclusions relies heavily on the proper specification of the statistical model, particularly concerning how the experimental factors relate to one another.

In a standard, or fixed-effects, ANOVA, all factors are typically assumed to be fully **crossed**. This means that every level of Factor A occurs in combination with every level of Factor B. For example, if we test three types of fertilizers (Factor A) and two levels of irrigation (Factor B), a fully crossed design would test all six possible combinations (Fertilizer 1/Irrigation Low, Fertilizer 1/Irrigation High, etc.). This approach is necessary to assess potential interactions between the factors, where the effect of one factor depends on the level of the other.

However, when the experimental setup precludes the crossing of factors--often due to logistical constraints, organizational hierarchy, or the nature of the sampling--a standard crossed ANOVA is inappropriate. If the treatments administered are unique to certain groups or settings, we introduce a nested structure. Recognizing this hierarchy is the first step toward applying the correct statistical model, ensuring that the error terms used for testing are appropriate for the specific source of variation being examined.

Defining the Nested Structure

A **nested ANOVA** is a specialized type of ANOVA in which at least one factor is *nested* inside another factor. This relationship implies that the levels of the nested factor are unique to, and therefore do not extend across, all levels of the main factor. In essence, the nested factor represents sub-units of analysis within the primary treatment groups, ensuring independence across primary factor levels.

Note: Sometimes a nested ANOVA is called a "hierarchical ANOVA." These two terms are often used interchangeably because they describe the hierarchical organization of the experimental units. The terminology emphasizes that the experimental design proceeds in stages, with lower-level factors being specific instances of higher-level factors.

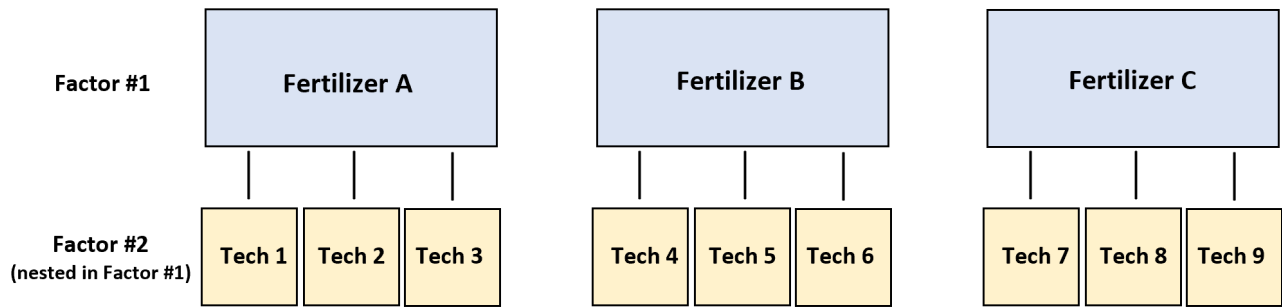
To illustrate the nested relationship, consider a multi-stage quality control process. If three different factories (Factor A) are each evaluated by three unique inspectors (Factor B), the factor "Inspector" is nested within the factor "Factory." Inspector 1 at Factory X is distinct from Inspector 1 at Factory Y; they are not the same person or entity being applied across all factories. If they were the same person inspecting all factories, the design would be crossed. Because the inspectors are unique to each factory, the factor "Inspector" is contained entirely within the levels of "Factory."

Practical Application: The Plant Growth Example

For a concrete example, suppose we would like to know if three different fertilizers (A, B, and C) produce different levels of plant growth. The response variable is **plant growth**, measured after a fixed period. The primary factor of interest is Fertilizer type, which has three levels.

To test this hypothesis while accounting for procedural variability, we implement a specific sampling design: we have three different technicians sprinkle fertilizer A on four plants each, another three technicians sprinkle fertilizer B on four plants each, and another three technicians sprinkle fertilizer C on four plants each. This design ensures that each technician works solely within one fertilizer group.

In this scenario, the dependent variable is plant growth and the two factors are Technician and Fertilizer. It turns out that Technician is *nested* within Fertilizer. The technician factor captures the variability in application method specific to each fertilizer group. We use this nesting to isolate variance components, recognizing that the technicians are distinct experimental units under each treatment level.



The resulting raw data reflects this hierarchical arrangement. Notice how the technician IDs (1, 2, 3) reset or restart within each major fertilizer group. This visualization confirms that the technicians are subordinate factors, specific only to the fertilizer they are assigned to handle. This distinction is crucial for correctly calculating the Sum of Squares and Degrees of Freedom.

Here's what the structure of the raw data would look like, demonstrating the unique assignment of the nested factor:

Plant Growth (inches)	Fertilizer	Technician
13	A	1
8	A	1
12	A	1
12	A	1
15	A	2
16	A	2
19	A	2
16	A	2
15	A	3
15	A	3
12	A	3
15	A	3
19	B	4
19	B	4
...

Testing Hypotheses with Nested ANOVA

When executing a nested ANOVA on this dataset (using statistical software like R, SPSS, or specialized packages), the model tests for the significance of the main factor (Fertilizer) and the nested factor (Technician within Fertilizer). The model is specifically designed to isolate the

variability introduced by the technicians' differing application methods from the true effect of the fertilizer itself. If the technician effect is significant, it indicates high within-group variability that must be accounted for.

In this scenario, a nested ANOVA can test for two primary effects:

Is plant growth significantly different across each level of factor 1 (**Fertilizer**)? This tests the main effect of the treatment, using the variability introduced by the nested factor as the error term.

Is plant growth significantly different across each level of factor 2 (**Technician** nested within Fertilizer)? This tests the variability introduced by the sub-units, which often serves as the appropriate error term for testing the main factor if the nested factor is considered random.

A key difference from a crossed design is the absence of an interaction term. Because technicians are not paired with all fertilizers, the interaction cannot be calculated. Instead, the variability due to the nested factor (Technician) is often used as the denominator in the F-ratio for the main factor (Fertilizer), assuming the nested factor is considered random. This adjustment ensures that the test for Fertilizer effect is appropriately conservative, accounting for the inherent variability introduced by the technicians within each group.

Interpreting the Nested ANOVA Output Table

When we perform a nested ANOVA, the output generated by statistical software will typically be presented in a summary table format similar to a standard ANOVA table, detailing the sources of variation, test statistics, and significance levels.

Source	Sum of Squares	df	Mean Square	F-Value	p-value
Fertilizer	372.7	2	186.3	53.238	< 0.0000
Technician(Fertilizer)	31.8	6	5.31	1.516	0.211
Residuals	94.5	27	3.5		

Here's how to interpret the crucial columns within the output:

Source: Identifies the component of variability being tested, including the primary factor (Fertilizer), the nested factor (Technician), and the residual error.

Sum of Squares: The total variation attributable to that specific source, calculated as the sum of the squared deviations from the appropriate mean.

df: The degrees of freedom, representing the number of independent pieces of information used to calculate the Sum of Squares. This value determines the shape of the F-distribution used for the test.

Mean Square: The estimated variance component for that source, calculated as the Sum of Squares divided by the corresponding degrees of freedom ($MS = SS / df$).

F-Value: The test statistic, calculated by dividing the Mean Square of the factor of interest by the appropriate Mean Square Error term. In the nested model, the F-Value for the main factor (Fertilizer) is often $MS(\text{Fertilizer}) / MS(\text{Technician})$.

p-value: The probability value that corresponds to the F-Value under the null hypothesis. This value is critical for making statistical decisions.

We analyze the p-value column to determine whether each factor exerts a statistically significant effect on plant growth, typically by comparing it against a predetermined significance level (alpha, usually 0.05). If the p-value is less than alpha, we reject the null hypothesis, concluding that the factor has a significant effect.

From the illustrative table above, we observe that fertilizer has a statistically significant effect on plant growth ($p\text{-value} < .05$), but technician (nested within Fertilizer) does not show a significant effect ($p\text{-value} = 0.211$).

This result provides valuable practical insight: if the goal is to maximize plant growth, researchers should focus on optimizing the fertilizer being used, as the variability introduced by the individual technician applying the fertilizer is statistically negligible in this experimental context. This reinforces the primary conclusion about the fixed factor.

Nested vs. Two-Way (Crossed) ANOVA

Understanding the distinction between nested and crossed designs is paramount for model selection. Choosing the wrong model--applying a standard Two-way ANOVA to a nested structure, or vice versa--will lead to miscalculation of the Sum of Squares, incorrect degrees of freedom, and ultimately, invalid statistical inferences about the factors.

Nested ANOVA's are distinct from two-way ANOVA's. In a nested ANOVA, at least one factor is nested *inside* another factor, meaning the levels of the subordinate factor are unique to the parent factor. This prevents the calculation of an interaction term, as not all combinations of factor levels are present.

This is fundamentally different from a Two-way ANOVA, which involves two or more factors where all factors are fully **crossed**. In a crossed design, every level of one factor is combined with every level of the other factor, allowing for the comprehensive assessment of main effects and their interaction effect (e.g., Fertilizer \times Technician).

For example, in the previous scenario suppose each technician sprinkled each type of fertilizer. In this case, we could perform a two-way ANOVA because every possible combination of technique

and fertilizer occurred in the dataset. If the researcher mistakenly used a nested model here, they would lose the ability to test the potentially significant interaction effect between the two factors, failing to fully characterize the experimental outcomes.

Further Considerations and Statistical Software

Here are a few important considerations regarding the application and flexibility of nested ANOVA designs:

Nested ANOVA's can accommodate more than two factors.

In the previous example the nested ANOVA had two factors, one nested inside the other. However, a nested ANOVA could have multiple levels of hierarchy, such as samples nested within plots, which are nested within sites, which are nested within regions. The complexity increases with each layer, and the selection of the correct error term for the F-test becomes highly dependent on the model structure.

Fixed vs. Random Effects.

Nested factors are frequently treated as **random effects**, meaning their specific levels are chosen randomly and are intended to represent a larger population of potential sub-units. If the nested factor is random, its associated Mean Square is used to test the significance of the fixed factor it is nested within. If all factors are fixed, the denominator for the F-test is typically the residual error, simplifying the interpretation but requiring careful consideration of the factor structure.

The implementation of a nested ANOVA requires specialized tools due to the complexity of calculating the correct error terms, especially in unbalanced designs. While basic spreadsheet software might handle simple, balanced cases, dedicated statistical packages are necessary for robust analysis and complex hierarchical data structures.

How to Perform a Nested ANOVA in Practice

The following tutorials explain how to perform a nested ANOVA using common statistical environments, emphasizing the importance of correctly defining the nesting structure within the software syntax:

Tutorial on specifying nested models in R (using the `aov()` or `lmer()` functions).

Step-by-step guidance for executing nested analysis using the General Linear Model (GLM) procedure in SPSS.