

How to Perform a Kruskal-Wallis One-Way ANOVA Test

Authored by
stats writer

January 22, 2026

RECOMMENDED CITATION

stats writer (2026). *How to Perform a Kruskal-Wallis One-Way ANOVA Test*.


PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=126996>

The Kruskal-Wallis One-Way ANOVA is a powerful non-parametric statistical test designed to assess whether there are statistically significant differences among three or more independent populations or groups. Serving as a crucial alternative to the standard parametric One-Way ANOVA, this method is employed primarily when the underlying assumptions of normality and homogeneity of variances are violated. Conceptually, it extends the principles of the Mann-Whitney U test to handle multiple group comparisons. The Kruskal-Wallis test operates by ranking all data points across all groups and then calculating a specialized test statistic (H-statistic). This procedure is especially valuable in research settings where data distributions are skewed or non-normal, allowing researchers to rigorously compare the medians of different groups. Named in honor of its developers, William Kruskal and Wilson Wallis, this test provides a robust framework for complex multi-group data analysis.

Defining the Kruskal-Wallis One-Way ANOVA

The **Kruskal-Wallis One-Way ANOVA** is specifically utilized in situations requiring the comparison of distribution centers (typically the median) across three or more distinct groups relative to a single variable of interest. This dependent variable must be **continuous**, meaning it can theoretically take on any value within a given range. A major advantage of using this non-parametric approach is its ability to handle variables that exhibit substantial **skewness** or deviation from a normal distribution, provided the spread of scores remains reasonably similar across the groups being analyzed.

Kruskal-Wallis One-Way ANOVA



This versatile statistical procedure is recognized by several alternative names in the literature, including the One-Way ANOVA on Ranks, Kruskal-Wallis One-Way Analysis of Variance, the **Kruskal-Wallis H Test**, or simply the Kruskal-Wallis Test. Researchers must ensure that the groups under investigation are **independent** (unrelated) and that the sample size provides adequate statistical power, typically requiring more than five data points in each group.

Key Assumptions Governing the Kruskal-Wallis H Test

Every valid statistical methodology relies on a set of core **assumptions** that the data must satisfy. These requirements ensure that the results derived from the statistical test are accurate, reliable, and interpretable. While the Kruskal-Wallis test is non-parametric and therefore less restrictive than parametric tests like the traditional ANOVA, it still requires careful adherence to specific criteria to prevent biased or erroneous conclusions.

The necessary assumptions for applying the Kruskal-Wallis One-Way ANOVA (2/5) are detailed below. It is important to confirm these conditions prior to running the analysis:

The dependent variable must be **Continuous**.

Data must originate from a **Random Sample**.

There must be **Enough Data** (Adequate Sample Size).

The distribution of scores must have a **Similar Spread Across Groups** (Homogeneity of Variance in Ranks).

The shapes of the distributions must be **Similar Between Groups**.

We will now explore the practical implications of each assumption individually.

Continuous Dependent Variable

The primary variable under investigation--the one you are measuring and comparing across your 3 or more groups--must be **continuous**. A continuous variable is defined as one that can assume any numerical value within a defined interval. This allows for fine-grained measurement and ranking, which is fundamental to how the Kruskal-Wallis procedure operates.

Excellent examples of continuous variables often encountered in research include metrics such as age (measured precisely), body weight, height, standardized test scores, comprehensive survey scores (if additive or averaged), and annual income. Variables that are strictly categorical or ordinal (like rankings or finishing places) are generally inappropriate for this test.

Requirement of a Simple Random Sample

For the results of the Kruskal-Wallis test to be generalizable to the broader population, the data points for each independent group must be derived from a simple **random sample**. This implies that every individual in the population had an equal chance of being selected for participation in the study. For instance, if a researcher aims to compare the effects of three different educational interventions, the participants receiving each intervention must be randomly selected or randomly assigned to ensure the groups are comparable at baseline.

The criticality of **random selection** cannot be overstated. When groups are not randomly determined, the resulting analysis may suffer from statistical **bias**, leading to incorrect or misleading conclusions about the population effects. Bias manifests as a systematic tendency to produce incorrect results due to poor data collection methodology.

*If the assumption of a random sample is violated, the conclusions drawn from the test are severely limited in scope. Researchers should always prioritize obtaining a simple random sample. Furthermore, if the data involves **paired samples** (i.e., multiple measurements taken from the same subjects across different conditions or time points), the Kruskal-Wallis test is inappropriate; in such cases, the Friedman Test should be utilized instead.*


Adequate Sample Size (Enough Data)

The test requires an adequate **sample size**, often referred to as having "enough data," to ensure the reliability of the calculated H-statistic. A general rule of thumb suggests that the sample size in each individual group should be greater than five observations. While some statistical texts might argue for slightly higher minimums, having more than five data points per group is usually sufficient for the asymptotic approximation of the chi-square distribution to be reliable.

It is crucial to note that the necessary sample size is also intrinsically linked to the **effect size**--the expected magnitude of the difference between groups. If researchers anticipate a large difference across the groups, a smaller sample size may suffice. Conversely, if the expected difference is subtle or small, a significantly larger sample size will be required to achieve sufficient statistical power to detect that difference reliably.

Similar Spread Across Groups (Homogeneity of Variance in Ranks)

This assumption requires that the variables exhibit a reasonably similar distribution of scores across all the independent groups. In the context of parametric statistics, this is known as **homogeneity of variance**. For the Kruskal-Wallis test, which uses ranks, we are essentially looking for similarity in the spread of the ranks across groups. To visually inspect this critical assumption, researchers should create appropriate plots of their data, such as box plots or dot plots, to assess whether the range and concentration of scores appear roughly equivalent between the groups being compared.



Group 2 is much more spread out than Group 1.



Group 1 and Group 2 have a similar spread

If your groups exhibit a substantially different spread or scale on the variable of interest--a clear violation of this assumption--the standard Kruskal-Wallis results may be misleading regarding differences in medians. In such instances, if the data were approximately normal, an alternative such as the Welch test ANOVA variant might be considered, although alternatives for non-parametric data with unequal variances are complex and context-dependent.

Similar Distribution Shape Between Groups

The interpretation of a significant Kruskal-Wallis result hinges on the assumption that the underlying population distributions of the groups are **similarly shaped** when visualized (e.g., as histograms or density plots). When the distribution shapes are similar, a significant result unambiguously indicates that the **medians** (or average ranks) of the groups are different. This allows for a clean comparison of central tendency.

If, however, the group distributions are highly dissimilar in shape (e.g., one is heavily positively

skewed and another is bimodal), then a significant test result indicates only that the distributions are different in some way--it could be due to differences in median, differences in spread, or differences in shape. If the shapes are dissimilar, researchers can confidently state that the overall distributions vary, but they cannot definitively argue that the difference is solely attributable to a difference in the average value or median.

Optimal Scenarios for Applying the Kruskal-Wallis Test

The decision to utilize the Kruskal-Wallis One-Way ANOVA (3/5) should be based on a clear set of criteria related to the research question and the nature of the data collected. This test provides a robust solution when the conditions for parametric testing are not met, yet a comparison of multiple independent groups is required.

The goal is to determine if many groups are **statistically different** on the variable of interest.

The dependent variable is **continuous**.

The study involves **3 or more groups**.

The samples are **independent** (unrelated).

The variable of interest is **skewed** or non-normally distributed.

A more detailed examination of these requirements will help clarify when the Kruskal-Wallis test is the most appropriate analytical tool.

Focus on Group Differences

The fundamental purpose of the Kruskal-Wallis test is to address a **difference question**: specifically, whether three or more predefined groups vary significantly concerning the measured outcome. This contrasts sharply with other statistical approaches, such as **correlation**, which examines the linear relationship between two variables, or **regression/prediction**, which attempts to model one variable based on the values of another. If the core research objective is simply to establish group disparity, the Kruskal-Wallis test provides the necessary inferential structure.

The Nature of the Dependent Variable

As previously established, the variable being measured must be **continuous**. This category encompasses data capable of fine-grained measurement, such as physiological metrics (e.g., heart rate), physical characteristics (e.g., height, weight), or performance metrics (e.g., speed or time to completion). The continuous nature of the data is essential because the Kruskal-Wallis procedure relies on converting these exact measurements into ranks before analysis.

It is critical to distinguish continuous data from other data types that are inappropriate for this test. These include **ordered or ordinal data** (like survey ratings, finishing places in a competition, or

business rankings), **categorical data** (such as gender, nationality, or eye color), or **binary data** (like 'yes/no' outcomes, disease presence/absence, or product purchase status). Use of non-continuous data types typically requires alternative non-parametric methods.

Comparison of Multiple Groups

The Kruskal-Wallis One-Way ANOVA is specifically designed for scenarios involving the comparison of **three or more independent groups** on a continuous dependent variable. It serves as the analytical bridge when traditional parametric methods are unsuitable due to data characteristics.

*If the study design involves only **two groups**, the appropriate non-parametric test is the Mann-Whitney U Test (2/5). Conversely, if the design involves only **one group** and the objective is to compare that group's distribution center to a hypothesized population parameter, the Single Sample Wilcoxon Signed-Rank Test should be employed.*

Requirement of Independent Samples

The Kruskal-Wallis test is strictly intended for **independent samples**, meaning that the observations in one group must be entirely unrelated to the observations in any other group. For instance, if researchers randomly select samples of individuals from three different geographic regions and measure their average daily screen time, the three regional groups are independent because no individual belongs to more than one group.

*In contrast, a **paired data** scenario--where the same subjects are measured repeatedly (e.g., pre-test, post-test, and follow-up)--violates the independence assumption. When using dependent or paired samples with three or more measurement points, the analysis must shift to the Friedman Test, which is the non-parametric equivalent of a repeated-measures ANOVA.*

Handling Non-Normal Data Distributions

One of the primary strengths of the Kruskal-Wallis One-Way ANOVA is its robustness against violations of the normality assumption. Unlike parametric tests which require the data to be approximately **normally distributed** (shaped like a bell curve), the Kruskal-Wallis test can be reliably performed even when the dependent variable is **highly skewed**. By utilizing ranks instead of raw scores, the test effectively minimizes the impact of extreme outliers and non-normal distribution shapes, making it a highly valuable tool for real-world research data.

A Practical Example of the Kruskal-Wallis Test

Consider a clinical trial designed to evaluate the efficacy of two novel medical treatments

compared to a control condition. This scenario perfectly illustrates the appropriate application of the Kruskal-Wallis test:

Group 1: Patients who received **Medical Treatment #1**.

Group 2: Patients who received **Medical Treatment #2**.

Group 3: Patients who received a **placebo** or standard control condition.

Variable of interest: Time in days required for the patient to fully recover from a disease.

In this experimental design, we have three clearly defined, **independent groups** and a single, **continuous variable** (time to recovery). Initially, a standard One-Way ANOVA might be considered. However, upon investigating the data, the researchers discover that the variable of interest, "time to recover," is heavily **positively skewed** (meaning most patients recovered quickly, but a few took a very long time). Due to this violation of the normality assumption, the researchers correctly elect to use the Kruskal-Wallis One-Way ANOVA (4/5) to compare the recovery times across the three intervention groups.

The foundational premise of the statistical test is framed by the **Null Hypothesis** (1/5). In statistical terminology, the Null Hypothesis asserts that there is no true effect of the treatments; specifically, it posits that the median recovery times for all three groups are identical. The study's objective is to determine if the data provide sufficient evidence to reject this Null Hypothesis, thereby confirming that at least one of the medical treatments (or the placebo) results in a significantly different average recovery time.

Once the experiment concludes, the Kruskal-Wallis analysis is performed, yielding two key statistical outputs: the **chi-square statistic** (or H-statistic) and the **p-value**. The chi-square statistic is a quantitative measure reflecting the overall magnitude of the differences observed among the three groups on the recovery-time variable. A larger chi-square value generally implies greater disparity between the groups' median ranks.

The **p-value** represents the probability of observing the current results (or results more extreme) if the Null Hypothesis (2/5) were true--i.e., if neither treatment had any actual effect on recovery time. Conventional statistical practice dictates that if the p-value is less than or equal to 0.05, the result is deemed **statistically significant**. This significance allows the researchers to reject the Null Hypothesis and confidently conclude that the observed difference is highly unlikely to be due to random chance alone.

A combination of a high chi-square statistic and a low p-value signifies that recovery time was significantly different in at least one group compared to the others. However, the Kruskal-Wallis test is an **omnibus test**; it only confirms that a difference exists somewhere among the groups. Further, targeted **post-hoc testing** (such as Dunn's test with Bonferroni correction) is essential to pinpoint exactly which specific group(s) had significantly higher or lower recovery times than the

rest.

ARABPSYCHOLOGY.COM