

How to Perform Mood's Median Test in R

Authored by
stats writer

December 26, 2025

RECOMMENDED CITATION

stats writer (2025). *How to Perform Mood's Median Test in R*. PSYCHOLOGICAL SCALES.
Retrieved from <https://scales.arabpsychology.com/?p=109015>

The Mood's Median Test is a fundamental non-parametric test designed to rigorously assess whether the population median differs across two or more independent groups. As a powerful alternative to parametric methods like ANOVA when underlying distributional assumptions are not met, this test relies on classifying data points relative to the combined sample median. Implementation in R is most effectively achieved using the `median_test()` function from the coin package, which provides the necessary test statistic (Z or χ^2) and the corresponding p-value for hypothesis testing.

The Mood's Median Test, formally known as the Brown-Mood Median Test, is specifically employed to compare the medians of two or more independent groups. It is especially useful when dealing with data that are ordinal or continuous but severely violate the assumptions of normality or homogeneity of variances required by parametric tests. This test operates by converting continuous data into a categorical structure based on the grand median of the entire combined sample, thereby simplifying the comparison into a $2 \times k$ contingency table, where k is the number of groups being compared.

The foundational concept involves calculating the overall median across all observations, regardless of group membership. Subsequently, for each group, the test counts how many observations fall above this grand median and how many fall below it. If the medians of the groups are truly identical, we would expect the proportion of observations above and below the grand median to be similar across all groups. A significant deviation from this expectation suggests that at least one group median differs from the others.

In the R environment, the execution of this analysis is facilitated by the `median_test` function available in the `coin` library. This function uses a formula syntax common in R statistical modeling, providing an intuitive way to specify the relationship between the response variable and the grouping variable:

```
median_test(response ~ group, data)
```

Understanding the Assumptions and Hypotheses

Before applying Mood's Median Test, it is vital to confirm that the underlying data structure meets the test's requirements. Primarily, the test requires that the samples are **independent**. This means that observations in one group must not influence or be related to observations in any other group. Furthermore, the dependent variable--the response variable being measured--must be measured on at least an ordinal scale, though it is most commonly applied to continuous data where the assumption of normality is suspect. If the data is truly nominal, this test would be inappropriate.

The formal statements for the hypotheses guiding Mood's Median Test are straightforward. The

goal is to compare the location parameters (the medians) across k independent populations. Let $\tilde{\mu}_i$ represent the population median for the i -th group:

Null Hypothesis (H_0): All population medians are equal ($\tilde{\mu}_1 = \tilde{\mu}_2 = \dots = \tilde{\mu}_k$). We assume no difference in central tendency across the groups.

Alternative Hypothesis (H_A): At least one population median is different from the others.

It is important to recognize that while this test compares medians, it specifically assesses whether the groups are different in location relative to the combined sample median. If the data distributions are highly dissimilar in shape or spread (heteroscedasticity), rejection of the null hypothesis might indicate differences in distribution rather than just differences in location. However, in most practical applications where the distributions are assumed to have similar shapes, the test provides robust evidence regarding the equality of the true population medians, offering a powerful non-parametric test alternative to ANOVA.

Setting Up the Analysis in R

Implementing Mood's Median Test in R requires loading the correct library and structuring the data appropriately. The `coin` package is essential, as it provides statistical tests based on conditional inference, resulting in highly accurate p-values, particularly for permutation-based non-parametric tests. Once the package is installed and loaded, the primary function used is `median_test()`.

The formula interface of `median_test()` requires that you identify the response and grouping variables within a specified data frame. A typical call looks like this:

```
median_test(response ~ group, data = my_dataframe)
```

This structure requires defining the relationship between the measured outcome and the grouping factor. Here is a breakdown of the variables specified in the function call:

response: This must be the **numeric vector** containing the outcome or response values (the dependent variable).

group: This must be the **factor or grouping vector** indicating which independent sample each response value belongs to (the independent variable).

data: This specifies the **data frame** containing both the response and group vectors, ensuring R can correctly interpret the context of the formula.

By using this formula interface, the `median_test()` function automatically handles the creation of the necessary two-by- k contingency table (where k is the number of groups) by identifying the grand median and counting observations above and below it for each group. This streamlined process eliminates the need for manual categorization, simplifying the analytical workflow significantly.

Example Scenario: Comparing Study Methods

To demonstrate the practical application of Mood's Median Test, consider a realistic scenario involving educational research. Suppose a teacher is investigating the efficacy of two distinct studying methods--Method 1 versus Method 2. The teacher hypothesizes that one method might yield significantly different exam scores compared to the other, but is cautious about using a standard t-test due to the small sample size and potential skewness in student performance data.

To test this robustly, the teacher randomly assigns 10 students to use Method 1 and another 10 students to use Method 2. After a two-week period, all 20 students take the same standardized exam. The resulting exam scores represent the response variable, and the study method represents the grouping variable. The teacher chooses Mood's Median Test because it is a non-parametric test that provides a reliable comparison of the central tendency (the median score) without requiring the assumption of normally distributed scores, thus ensuring the validity of the conclusion even with limited data.

The analysis aims to determine if there is a statistically significant difference in the central exam performance--as measured by the median--between students using Method 1 and those using Method 2. We will now proceed to the data preparation and execution steps in R.

Step 1: Creating and Viewing the Data Frame

The first crucial step in any R analysis is properly structuring the data into a **data frame**. We define the grouping vector (`method`) and the continuous response vector (`score`). Note that the `method` variable is defined as a factor by repeating each level ten times, corresponding to the twenty individual scores recorded.

We begin by defining the two vectors and combining them into a data frame named `examData`. The code chunk below shows the creation and structure of this data set, illustrating how R displays the grouped observations, ensuring the data is correctly set up for the `median_test()` function.

```
#create data
method = rep(c('method1', 'method2'), each=10)
score = c(75, 77, 78, 83, 83, 85, 89, 90, 91, 97, 77, 80, 84, 84, 85, 90, 92, 92, 94, 95)
examData = data.frame(method, score)

#view data
examData

method score
1 method1 75
```

```
2 method1 77
3 method1 78
4 method1 83
5 method1 83
6 method1 85
7 method1 89
8 method1 90
9 method1 91
10 method1 97
11 method2 77
12 method2 80
13 method2 84
14 method2 84
15 method2 85
16 method2 90
17 method2 92
18 method2 92
19 method2 94
20 method2 95
```

This data frame clearly shows 20 observations, equally split between the two studying methods. This structure ensures that R can correctly identify the independent (`method`) and dependent (`score`) variables required for the statistical test in the subsequent step. Proper data preparation is critical for avoiding errors when calling the analytical function.

Step 2: Performing Mood's Median Test and Interpreting the Results

With the data frame correctly prepared, we proceed to perform the core analysis by loading the necessary `coin` library and then applying the `median_test()` function using the standard formula notation, `score ~ method`. This tells R to compare the distribution of `score` across the levels of the `method` variable.

The output provided by `median_test()` is based on the asymptotic distribution of the test statistic, which approximates the Chi-squared distribution used for calculating the p-value. The following code executes the test and displays the results:

```
#load the coin library
```

```
library(coin)
```

```
#perform Mood's Median Test
```

```
median_test(score~method, data = examData)

#output
Asymptotic Two-Sample Brown-Mood Median Test

data: score by method (method1, method2)
Z = -0.43809, p-value = 0.6613
alternative hypothesis: true mu is not equal to 0
```

The result yields a Z-statistic of -0.43809 and, crucially, a p-value of **0.6613**. Since this p-value (0.6613) is significantly larger than the standard significance level of 0.05, we **fail to reject the null hypothesis**. This means we do not possess sufficient evidence to conclude that there is a statistically significant difference in the population median exam scores between the students who used Method 1 and those who used Method 2.

Advanced Usage: Adjusting for Ties with `mid.score``

A statistical subtlety in Mood's Median Test arises when observations are exactly equal to the grand sample median (known as ties). By default, the `median_test()` function in the `coin` package assigns a score of 0 to these tied observations, effectively excluding them from the 'above' or 'below' counts, which is a conservative approach. However, researchers might prefer a less conservative test, which is achieved by adjusting the `mid.score` argument.

The `mid.score` argument controls how these tied observations contribute to the counts. A common alternative is setting `mid.score="0.5"`. When this setting is used, observations exactly equal to the median are treated as contributing 0.5 to the count of observations "above the median" and 0.5 to the count of observations "below the median" for their respective groups. This approach often leads to a slightly more powerful test by utilizing more information about the data distribution near the center.

For example, the following code performs the exact same Mood's Median Test but incorporates the `mid.score="0.5"` adjustment to handle any ties symmetrically:

```
#perform Mood's Median Test with mid.score adjustment
median_test(score~method, mid.score="0.5", data = examData)

#output
Asymptotic Two-Sample Brown-Mood Median Test

data: score by method (method1, method2)
Z = -0.45947, p-value = 0.6459
```

alternative hypothesis: true mu is not equal to 00

As demonstrated, the adjustment changes the Z statistic and the resulting p-value (from 0.6613 to 0.6459). While the ultimate conclusion in this specific example remains unchanged (we still fail to reject H_0), this feature is crucial for analysts working with small sample sizes or data sets where ties around the median are frequent, allowing for careful consideration of the statistical power and interpretation.

Choosing Mood's Median Test vs. Other Non-Parametric Options

While Mood's Median Test is straightforward and robust, analysts should be aware of its relative statistical power. Compared to other location-based non-parametric tests, such as the Wilcoxon Rank-Sum Test (for two groups) or the Kruskal-Wallis H Test (for multiple groups), Mood's Median Test is often considered less powerful. This difference stems from the fact that Mood's test only utilizes the information of whether an observation is above or below the combined median, ignoring the actual magnitude of the rank difference.

The Kruskal-Wallis test, by contrast, uses the ranks of all observations across all groups, thereby incorporating more detail about the relative scores, which generally leads to greater power to detect differences. Therefore, if the primary concern is maximizing statistical power and the data is continuous, Kruskal-Wallis is often the preferred non-parametric test.

However, Mood's Median Test retains significant value in specific contexts. It is the ideal choice when the researcher is strictly interested in comparing the central location in terms of the median, or when the data is strongly affected by outliers, making a highly robust method necessary. Furthermore, due to its simplicity in transforming the data into a binary outcome, it provides results that are exceptionally easy to interpret and communicate, making it a reliable tool for initial exploratory analysis in R.