

Learn How to Calculate Trimmed Mean in R with Examples

Authored by
Mohammed loot

November 2, 2025

RECOMMENDED CITATION

Mohammed loot (2025). *Learn How to Calculate Trimmed Mean in R with Examples*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=8695>

A **trimmed mean**, sometimes referred to as a truncated **mean**, is a fundamental measure of central tendency used in statistical analysis. Unlike the standard arithmetic mean, the trimmed mean is calculated by systematically removing a specified percentage of the smallest and largest values from a given **dataset**. This process enhances the robustness of the average calculation.

For instance, calculating a 10% trimmed mean means that the top 10% of the highest values and the bottom 10% of the lowest values are discarded before the arithmetic average is computed. This technique is particularly valuable when dealing with distributions that are prone to **outliers** or extreme observations that could otherwise severely skew the reported average.

In the statistical programming environment **R**, calculating this specialized average is straightforward, relying on a simple argument within the built-in `mean()` function. The following sections provide a detailed walkthrough of the R syntax and practical examples demonstrating how to apply this calculation to various data structures.

Understanding the Trimmed Mean and Robust Statistics

The primary motivation for using a **trimmed mean** is to achieve greater resilience against influential data points. When a dataset contains outliers--values that lie abnormally far from other values--the traditional arithmetic mean can be significantly biased, potentially leading to misleading conclusions about the typical value of the distribution.

By removing these extreme observations from both tails of the distribution, the trimmed mean provides a more accurate representation of the center of the majority of the data. The percentage of trimming chosen (e.g., 5%, 10%, or 20%) directly reflects the degree of robustness desired; a higher percentage removes more potential outliers but also sacrifices more valid data points.

It is important to note the difference between a trimmed mean and a median. The **median** is essentially a 50% trimmed mean, as it eliminates all values except the single middle observation (or the average of the two middle observations). The trimmed mean serves as a useful compromise between the sensitivity of the standard mean (0% trim) and the extreme robustness of the median.

Implementing the Trimmed Mean using R Syntax

The statistical software **R** simplifies the computation of the trimmed mean through the versatile `mean()` function. Instead of manually sorting and subsetting the data, we only need to utilize the dedicated `trim` argument, which accepts a value between 0 and 0.5.

The value specified for `trim` represents the fraction of observations to be removed from **each end** of the ordered data. Therefore, if you specify `trim=0.1`, a total of 20% of the data (10% smallest

and 10% largest) will be excluded from the final [mean](#) calculation.

The basic syntax structure below demonstrates how to calculate a 10% trimmed mean for a generic data object `x`. This simple approach is applied consistently whether `x` is a simple vector or a column extracted from a larger [data frame](#).

```
# Calculate the 10% trimmed mean for object x  
mean(x, trim=0.1)
```

The following practical examples demonstrate the application of this function across various common scenarios encountered during data analysis in R.

Example 1: Calculating the Trimmed Mean for an R Vector

The most straightforward application of the trimmed mean function is on a simple numeric [vector](#). A vector is the most basic data structure in [R](#), typically representing a single sequence of measurements or observations.

In this scenario, we define a vector containing 20 raw observations. We then apply the `mean()` function with the `trim=0.1` argument to calculate the 10% trimmed mean. Since the [dataset](#) contains 20 values, a 10% trim means 2 values will be removed in total (1 smallest, 1 largest).

```
# Define the sample data vector  
data = c(22, 25, 29, 11, 14, 18, 13, 13, 17, 11, 8, 8, 7, 12, 15, 6, 8, 7, 9, 12)  
  
# Calculate the 10% trimmed mean  
mean(data, trim=0.1)
```

```
12.375
```

The resulting 10% trimmed [mean](#) is precisely **12.375**. If the standard (untrimmed) mean were calculated, the presence of the extreme values (6 and 29) might slightly shift the result. By applying the trim, we ensure the average reflects the central distribution of the remaining 18 observations.

Example 2: Applying the Trimmed Mean to a Data Frame Column

In practical data science, data is often organized into [data frames](#), which function much like tables or spreadsheets. To calculate the trimmed mean for a specific variable, we must first select the desired column using the dollar sign (`$`) notation.

In the following demonstration, we create a data frame containing statistics for several fictional

players (points, assists, and rebounds). We are interested in calculating the 5% trimmed mean specifically for the `points` column. Since 5% of 8 observations is 0.4, R rounds this calculation, effectively removing the single smallest and single largest value from the set of 8 points before calculating the average.

Create the sample data frame

```
df = data.frame(points=c(25, 12, 15, 14, 19, 23, 25, 29),
  assists=c(5, 7, 7, 9, 12, 9, 9, 4),
  rebounds=c(11, 8, 10, 6, 6, 5, 9, 12))
```

```
# Calculate the 5% trimmed mean of the 'points' variable
mean(df$points, trim=0.05)
```

```
20.25
```

The calculated 5% trimmed mean for the values in the `points` column is **20.25**. This measure represents the average scoring rate after the two most extreme values (the highest and the lowest point totals) have been excluded from the computation, thereby providing a more central estimate.

Example 3: Calculating Trimmed Means Across Multiple Columns

A common requirement in data analysis is calculating a summary statistic, such as the trimmed mean, for several variables simultaneously within the same [data frame](#). While looping structures could be used, the R base function [sapply\(\)](#) offers a concise and vectorized way to achieve this efficiency.

The [sapply\(\)](#) function applies a specified function (in this case, the `mean()` function with `trim=0.05`) across a list or a subset of columns. We provide `sapply()` with a subset of the data frame `df`, selecting both the `points` and `assists` columns for processing.

Create the sample data frame (redefined for clarity)

```
df = data.frame(points=c(25, 12, 15, 14, 19, 23, 25, 29),
  assists=c(5, 7, 7, 9, 12, 9, 9, 4),
  rebounds=c(11, 8, 10, 6, 6, 5, 9, 12))
```

```
# Calculate 5% trimmed mean for both 'points' and 'assists' columns
sapply(df, function(x) mean(x, trim=0.05))
```

```
points assists
20.25 7.75
```

The output from the `sapply()` function clearly presents the trimmed means for all selected columns. This method is highly scalable and efficient when working with [R datasets](#) containing dozens or hundreds of variables.

Based on the results, we can quickly summarize the robust central tendency for each metric:

The 5% trimmed [mean](#) of the `points` column is **20.25**.

The 5% trimmed mean of the `assists` column is **7.75**.

Considerations for Choosing the Trimming Percentage

Selecting the appropriate percentage for the `trim` argument is a critical step that depends heavily on the underlying data distribution and the specific goals of the analysis. There is no universally correct percentage; the choice often involves a trade-off between bias reduction and variance inflation.

If the [dataset](#) is suspected of containing severe, distant outliers (e.g., recording errors or genuinely exceptional events), a higher trim value (like 10% or 20%) is often recommended to ensure these points do not distort the central estimate. This yields a more **robust estimator**.

Conversely, if the data is known to be relatively clean or symmetrically distributed without extreme values, a lower trim value (such as 0.05) or even the standard mean (`trim=0`) is appropriate. Over-trimming data that lacks outliers can unnecessarily discard valuable information, leading to increased variance in the resulting [mean](#). Analysts should always visualize their data distribution before settling on a final trimming percentage.

Related: For users interested in exploring other methods for handling skewed distributions and outliers, tutorials on calculating the median or Winsorized mean may also be beneficial.

Additional Resources

The following tutorials provide additional information about trimmed means and related robust statistical methods in [R](#):