

Learning to Apply Functions to Rows in R with dplyr

Authored by
Mohammed loot

October 29, 2025

RECOMMENDED CITATION

Mohammed loot (2025). *Learning to Apply Functions to Rows in R with dplyr*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=5186>

In the vast ecosystem of [R](#) programming, especially when dealing with structured, tabular datasets, one task consistently challenges developers: applying a custom [function](#) across each row of a [data frame](#). While base R offers methods to accomplish this, the process often lacks the clarity and efficiency modern data science demands. Fortunately, the widely adopted [dplyr](#) package, a foundational component of the [Tidyverse](#), offers an elegant, powerful, and highly readable solution. This comprehensive guide details how to leverage [dplyr](#) to execute row-wise function applications, significantly streamlining your data manipulation workflows and ensuring code that is both performant and easy to maintain.

The Challenge of Row-Wise Operations in R

Data analysis frequently requires calculations that traverse horizontally, focusing on individual observations or records rather than summarizing vertical columns. Consider scenarios where you need to calculate a composite score for a single entity (a student, a sensor, or a company) based on values spread across several columns. For example, finding the highest sales figure recorded in any of the four quarterly columns, or calculating the average response time for a single user across five different trials. These are inherently **row-wise operations**.

Historically, achieving these calculations in base [R](#) could be cumbersome. Solutions often involved writing explicit loops (which can be slow and less R-idiomatic) or utilizing the multifaceted `apply()` family of functions, which, while powerful, can sometimes lead to complex and less intuitive code, particularly for data analysts new to the language. The core difficulty lay in efficiently instructing R to treat the values within a single row as a distinct collection upon which a function should operate.

This is precisely where [dplyr](#) excels. Designed to provide a consistent and expressive set of data manipulation verbs, [dplyr](#) simplifies complex transformations. For operations requiring row-level granularity, the package introduces the pivotal `rowwise()` function. By setting the context for subsequent operations, `rowwise()` allows functions like `mutate()` to apply calculations independently to every row, transforming a challenging horizontal task into a straightforward, chained sequence of commands.

Core Syntax for Row-Wise Function Application with dplyr

The standard methodology for applying a function to each row of a [data frame](#) using [dplyr](#) revolves around the seamless chaining of operations, made possible by the [pipe operator](#) (`%>%`). This operator elegantly passes the output of one step directly as the primary input to the next, creating a clear, left-to-right flow that mirrors human thought processes in data transformation.

The foundational pattern for any row-wise operation is structured as follows:

```
df %>%
```

```
rowwise() %>%  
mutate(mean_value = mean(c(col1, col2, col3), na.rm=TRUE))
```

Let's dissect this sequence to understand its power. The process starts with `df`, representing your initial data structure. This is immediately piped into the `rowwise()` function, which is the crucial step: it effectively converts the [data frame](#) into a specialized "row-grouped" [tibble](#), ensuring that subsequent functions are applied one row at a time. The next step utilizes `mutate()`, the workhorse function for creating new columns. Here, we define a new column, `mean_value`, whose values are calculated by applying the `mean()` function to the vector created by combining the values from `col1`, `col2`, and `col3` within that specific row. A vital inclusion is `na.rm=TRUE`, which dictates that any [missing values](#) (NA) found within the row are ignored during the calculation, preventing the result from automatically defaulting to NA.

The true advantage of this method lies in its adaptability. While the example above computes the mean, you are free to swap out `mean()` for nearly any other R function--be it `sum()`, `max()`, `min()`, `sd()` (standard deviation), or even a custom analytical [function](#) you define yourself. This versatility makes the `rowwise()` approach an indispensable technique for highly customized and specialized data analysis.

Setting Up Our Example Data

To provide a clear, practical demonstration of the `rowwise()` syntax in action, we will construct a hypothetical dataset. This [data frame](#) models points scored by several basketball players across three separate games. This realistic scenario allows us to showcase how `dplyr` handles row-wise calculations, including the common challenge of dealing with incomplete data.

We generate our example [data frame](#) using the following base R code:

```
#create data frame  
df <- data.frame(game1=c(22, 25, 29, 13, 22, 30),  
game2=c(12, 10, 6, 6, 8, 11),  
game3=c(NA, 15, 15, 18, 22, 13))
```

```
#view data frame  
df
```

```
game1 game2 game3  
1 22 12 NA  
2 25 10 15  
3 29 6 15  
4 13 6 18
```

```
5 22 8 22
6 30 11 13
```

The resulting `df` contains score data organized into columns `game1`, `game2`, and `game3`, with each row representing a unique player's performance. Importantly, the score for the first player in `game3` is recorded as [NA](#) (Not Applicable). This intentionally introduced [missing value](#) will serve as a critical test case, allowing us to confirm that the `na.rm=TRUE` argument correctly handles incomplete data without invalidating our row-wise calculations.

Example 1: Calculating the Mean of Specific Columns Per Row

Our first analytical goal is to compute the average points scored for each player across two specific games: `game1` and `game3`. This operation is highly relevant in performance analysis, providing a consolidated metric for individual entities based on a subset of data points.

The following code utilizes [dplyr](#) to execute this calculation:

```
library(dplyr)
```

```
#calculate mean of game1 and game3
df %>%
  rowwise() %>%
  mutate(mean_points = mean(c(game1, game3), na.rm=TRUE))
```

```
# A tibble: 6 x 4
```

```
# Rowwise:
```

```
game1 game2 game3 mean_points
```

```
1 22 12 NA 22
```

```
2 25 10 15 20
```

```
3 29 6 15 22
```

```
4 13 6 18 15.5
```

```
5 22 8 22 22
```

```
6 30 11 13 21.5
```

We begin by loading the necessary [dplyr](#) library. The data frame `df` is then piped to `rowwise()`, setting the stage for row-specific calculations. Subsequently, `mutate()` generates the new column, `mean_points`. Within `mutate()`, the `mean()` function calculates the average of the values combined from `game1` and `game3` for the current row. The inclusion of `na.rm=TRUE` is critical, as it ensures that the [missing value](#) in the first row is excluded from the denominator, resulting in a

correct calculation.

By reviewing the output, we can confirm the appropriate handling of data, especially in the presence of missing scores:

For the first player, `game1` is **22** and `game3` is [NA](#). Because `na.rm=TRUE` is applied, only 22 is used, yielding a `mean_points` of **22**.

The second player scored **25** and **15**. The mean is calculated as $(25 + 15) / 2 = 20$.

The third player scored **29** and **15**, resulting in a mean of **22**.

This demonstrates the primary advantage of `rowwise()` : enabling concise and accurate row-level aggregation, even when dealing with sparse data.

Example 2: Determining the Maximum Value Per Row

In many analytical contexts, the focus shifts from the average to the peak performance or maximum recorded value. This is useful for identifying outliers, maximum load, or highest scores achieved. For our basketball data, we will determine the maximum points scored by each player across the pair of games, `game2` and `game3`.

To perform this peak calculation, we simply substitute the `mean()` function with the `max()` function inside the `mutate()` step:

library(dplyr)

```
#calculate max of game2 and game3
df %>%
  rowwise() %>%
  mutate(max_points = max(c(game2, game3), na.rm=TRUE))
```

```
# A tibble: 6 x 4
```

```
# Rowwise:
```

```
game1 game2 game3 max_points
```

```
1 22 12 NA 12
```

```
2 25 10 15 15
```

```
3 29 6 15 15
```

```
4 13 6 18 18
```

```
5 22 8 22 22
```

```
6 30 11 13 13
```

The pattern remains consistent: `rowwise()` prepares the [data frame](#) for row-level processing, and

`mutate()` introduces the `max_points` column. By employing the `max()` function, we efficiently identify the largest score between `game2` and `game3` for every player. Once again, `na.rm=TRUE` ensures that the calculation proceeds correctly, even when a score is marked as [NA](#).

Examination of the generated `max_points` column confirms the maximum value calculation:

For the first player (Row 1), comparing **12** (game2) and [NA](#) (game3) results in a maximum of **12**.

The second player (Row 2) scored **10** and **15**, yielding a maximum of **15**.

The fourth player (Row 4) scored **6** and **18**, resulting in a maximum of **18**.

This simple replacement of the function within `mutate()` showcases the flexibility of the `rowwise()` method, allowing users to rapidly switch between different row-wise aggregate statistics.

Example 3: Calculating Standard Deviation Across Rows

Beyond simple averages and maximums, advanced statistical measures are often required. The [standard deviation](#) (SD) is a crucial metric that reveals the consistency or volatility of data points. Calculating the [standard deviation](#) across rows for specific columns allows us to assess the variance in performance for each player.

To calculate the [standard deviation](#) of points between `game2` and `game3` for every player, we use the `sd()` function:

```
library(dplyr)
```

```
#calculate standard deviation of game2 and game3
df %>%
  rowwise\(\) %>%
  mutate(sd_points = sd(c(game2, game3), na.rm=TRUE))
```

```
# A tibble: 6 x 4
```

```
# Rowwise:
```

```
game1 game2 game3 sd_points
```

```
1 22 12 NA NA
```

```
2 25 10 15 3.54
```

```
3 29 6 15 6.36
```

```
4 13 6 18 8.49
```

```
5 22 8 22 9.90
```

```
6 30 11 13 1.41
```

The resulting output reveals important statistical considerations regarding single-point calculations. While `na.rm=TRUE` successfully excludes the [missing value](#) in Row 1, the `sd()` function requires at least two valid data points to calculate a meaningful measure of dispersion. Therefore, for the first row, where only the score of **12** remains, the `sd_points` result is also **NA**.

Conversely, for players with two scores, the [standard deviation](#) reflects their consistency:

The second player (scores **10** and **15**) has an SD of approximately **3.54**, indicating moderate variability.

The third player (scores **6** and **15**) has a higher SD of **6.36**, suggesting a greater spread or less consistency between their two scores.

This example underscores a key principle of data analysis: while `rowwise()` provides the structure for calculation, the intrinsic mathematical requirements of the chosen [function](#) must always be respected for the results to be statistically valid.

Understanding the `rowwise()` Function in Detail

To fully master row-wise operations, it is essential to grasp the mechanism behind `rowwise()`. Its core purpose is to temporarily redefine the grouping structure of your data. By transforming the [data frame](#), it essentially treats every single row as a distinct group of size one. When [dplyr](#) verbs like `mutate()` or `summarise()` are applied to a grouped data structure, they execute the operation independently within each group. Thus, by making each row its own group, subsequent calculations are guaranteed to run row-by-row.

It is important to note that for simple operations across all columns, base [R](#) provides optimized alternatives such as `rowSums()` or `rowMeans()`. However, when dealing with selective columns (as shown in our examples) or when applying highly specialized, non-standard [functions](#), the combination of `rowwise()` and `mutate()` offers unmatched flexibility and clarity.

A crucial best practice following any `rowwise()` operation is the use of `ungroup()`. Since `rowwise()` imposes a persistent grouping attribute onto the data structure, this grouping will carry forward to subsequent steps in your data pipeline. If those later steps are designed to operate on an ungrouped structure (or grouped by different variables), unexpected errors or incorrect aggregated results may occur. By explicitly calling `ungroup()` after the row-wise calculation is complete, you reset the data back to its standard, ungrouped [tibble](#) form, ensuring predictable behavior for the rest of your script. For deeper technical insights, always refer to the [official documentation for dplyr](#).

Conclusion and Further Exploration

Applying custom or standard [R](#) functions to specific columns on a row-by-row basis is a fundamental capability in data analysis. By harnessing [dplyr](#)'s elegant `rowwise()` function in conjunction with `mutate()`, you gain a highly readable, efficient, and flexible method for performing these horizontal computations. We have successfully demonstrated this approach across various statistical measures--mean, maximum, and [standard deviation](#)--highlighting the importance of correctly handling [missing values](#).

The core syntax remains the powerful template for all your row-wise needs: **`df %>% rowwise() %>% mutate(new_col = your_function(c(col1, col2), na.rm=TRUE))`**. Remember that `your_function` is entirely interchangeable, offering the ability to implement virtually any calculation necessary for your analysis. Furthermore, always be mindful of the statistical requirements of the function you choose; as demonstrated with `sd()`, some functions require a minimum number of valid data points, even when missing data is removed.

Mastery of `rowwise()` simplifies complex data preparation tasks, making your [R](#) code cleaner and your analytical insights more reliable. We encourage extensive experimentation with diverse functions and datasets to solidify this valuable data manipulation skill.

Additional Resources

For those looking to expand their [dplyr](#) expertise, the following tutorials explain how to perform other common tasks using this powerful package: