

How to Easily Remove the Last Row from a Data Frame Using dplyr

Authored by
stats writer

November 21, 2025

RECOMMENDED CITATION

stats writer (2025). *How to Easily Remove the Last Row from a Data Frame Using dplyr*. PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=98772>

Efficiently manipulating and cleaning data is a fundamental requirement in statistical computing, especially when working with large datasets in R. One common task involves removing specific observations, such as the final row, from a data frame. While several base R methods exist, the dplyr package, part of the Tidyverse ecosystem, offers a highly readable and powerful approach using the filter() function.

Initially, one might consider filtering based on a unique identifying column, such as finding the maximum value of an 'ID' column and excluding it. However, this approach is fundamentally unreliable. It assumes the last row always corresponds to the maximum ID value, which is not guaranteed after sorting or complex manipulations, and it fails entirely if the data frame lacks a sequential identifier. Furthermore, relying on column values rather than physical row position makes the code less robust and difficult to generalize. A far more stable and programmatic solution involves leveraging dplyr's built-in window functions designed specifically for positional row operations.

The core of this elegant solution lies in combining the filter() verb with two specific functions: row_number(), which identifies the current position of the row, and n(), which calculates the total count of rows within the context of the current operation. By setting up a logical comparison that excludes the highest row number, we can precisely target and excise the final observation(s) from the dataset, ensuring accurate removal regardless of the data frame's contents or size. The methods below demonstrate how to achieve this for both single and multiple rows.

The Power of the dplyr Package

The dplyr package is the go-to resource for data manipulation in R, providing a consistent set of verbs that greatly simplify complex data transformations. The primary advantage of using this package is the enhanced readability and expressiveness of the code, which makes scripts easier to maintain and debug. Unlike base R subsetting methods that often rely on bracket notation, dplyr uses function calls and the pipe operator (`%>%`) to chain operations logically.

To remove rows based on their position, we rely on the filter() function. This function retains all rows for which the logical expression evaluates to `TRUE`. To remove the last row, therefore, we must construct an expression that evaluates to `TRUE` for every row *except* the final one. This requires internal knowledge of the row count and the current row's index, capabilities which are provided seamlessly by dplyr's window functions.

Understanding the context of row operations is vital. When we use filter(), the expression is evaluated row by row. This is where row_number() and n() become essential, allowing us to perform relative comparisons. The combination ensures that the code remains concise while performing a sophisticated positional subtraction operation necessary for precise data subsetting.

Core Concepts: `row_number()` and `n()`

To successfully remove the last row of a [data frame](#) using [dplyr](#), a solid understanding of two key functions is necessary: `row_number()` and `n()`. These functions are part of [dplyr](#)'s suite of window functions, meaning they calculate values across a group of rows (in this case, the entire data frame) but are often applied within row-wise operations.

The function `n()` is extremely straightforward: it calculates and returns the total number of rows (or observations) in the current group being processed. When used outside of a `group_by()` operation, `n()` simply provides the total length of the [data frame](#). If a data frame has 100 rows, `n()` will evaluate to 100 in the context of the `filter()` call. This count serves as our reference point for identifying the endpoint of the dataset.

Conversely, `row_number()` assigns a sequential integer to each row, starting from 1 up to the total count provided by `n()`. This function is dynamic and calculates the position relative to the current ordering of the data. By combining these two functions, we can create a powerful logical condition: `row_number() <= n() - 1`. This expression instructs R to keep every row whose number is less than or equal to the total number of rows minus one, thus effectively excluding only the very last row, which corresponds precisely to the index `n()`.

Method 1: Removing the Single Last Row

To remove the single final observation from a [data frame](#), we apply the `filter()` function using the positional logic discussed above. This method is highly efficient and scalable, requiring only a single line of code after loading the necessary library. This approach is superior to manually calculating the row count and then subsetting using bracket notation in base R, as the entire operation is self-contained within the readable [dplyr](#) syntax.

The core syntax involves piping the data frame (`df`) into the `filter()` function. Inside the filter, we set the condition `row_number() <= n() - 1`. The subtraction of 1 from the total row count (`n()`) determines the cutoff point. If the data frame has 10 rows, `n() - 1` equals 9. The condition therefore ensures that only rows 1 through 9 are retained, discarding row 10.

Below is the concise code required to implement this technique. It is crucial to remember to load the `dplyr` package first, as these specialized functions are not available in base R. This approach guarantees that the last row, regardless of its underlying data values, is always excluded.

Method 1: Remove Last Row from Data Frame

```
library(dplyr)
```

```
#remove last row from data frame
```

```
df <- df %>% filter(row_number() <= n()-1)
```

Method 2: Removing the Last N Rows

Extending Method 1 to remove multiple rows is straightforward and demonstrates the flexibility of using `n()`. If the requirement is to truncate the dataset by removing the last N observations rather than just one, we simply adjust the subtraction factor in the filtering expression. This capability is highly useful in scenarios like time-series analysis where recent outliers or incomplete trailing entries must be discarded.

For instance, to remove the last four rows, the logical condition becomes `row_number() <= n() - 4`. If the original data frame contains 10 rows, `n() - 4` evaluates to 6. Consequently, only rows 1 through 6 are retained, successfully discarding rows 7, 8, 9, and 10. This proportional logic is robust and requires no external variables to manage the operation.

It is essential to ensure that the number subtracted (N) is less than the total number of rows. If N is greater than or equal to the total row count, the resulting data frame will be empty. This method provides a reliable way to truncate data dynamically without relying on absolute row indices, maintaining code clarity and efficiency.

Method 2: Remove Last N Rows from Data Frame

```
library(dplyr)
```

```
#remove last four rows from data frame
```

```
df <- df %>% filter(row_number() <= n()-4)
```

Understanding the Row Count Logic

The `n()` function extracts the total number of rows in the data frame or the current grouping context. This is the cornerstone of the positional filtering technique in `dplyr` when dealing with relative positions like "the last row." It provides the necessary numerical reference point from which to calculate the desired subset size.

By utilizing the condition `row_number() <= n() - N`, we are explicitly instructing the `filter()` function to subset the data frame. This condition ensures that we only retain rows where the sequential row number is less than or equal to the total number of rows (N) minus the count of rows we wish to discard. This elegant mathematical relationship handles all boundary conditions related to removing elements from the end of a vector or data structure.

To illustrate these methods practically, we will use a sample data frame representing team

statistics. This example will clearly show how the row indices are affected after each filtering operation, confirming that only the targeted final rows are removed.

#create data frame

```
df <- data.frame(team=c('A', 'A', 'A', 'B', 'B', 'C', 'C', 'C'),
points=c(18, 13, 19, 14, 24, 21, 20, 28),
assists=c(5, 7, 17, 9, 12, 9, 5, 12))
```

```
#view data frame
```

```
df
```

```
team points assists
```

```
1 A 18 5
```

```
2 A 13 7
```

```
3 A 19 17
```

```
4 B 14 9
```

```
5 B 24 12
```

```
6 C 21 9
```

```
7 C 20 5
```

```
8 C 28 12
```

Practical Implementation: Removing the Last Row (Example 1)

Using the sample data frame `df`, which contains eight observations, we will now apply Method 1 to remove the final row. Our target is row 8 (`c(28, 12)`). By executing the filter condition `row_number() <= n() - 1`, we instruct R to keep all rows whose index is less than or equal to 8 - 1, or 7.

The following code demonstrates the complete workflow, starting with loading the library and applying the filter operation directly to the existing data frame, overwriting the variable `df` with the truncated result. This is a common pattern in data wrangling to permanently update the dataset.

```
library(dplyr)
```

```
#remove last row from data frame
```

```
df <- df %>% filter(row_number() <= n()-1)
```

```
#view updated data frame
```

```
df
```

```
team points assists
```

```
1 A 18 5
2 A 13 7
3 A 19 17
4 B 14 9
5 B 24 12
6 C 21 9
7 C 20 5
```

As clearly shown in the output, the resulting data frame now contains only seven rows. The final observation from the original data (index 8) has been successfully and unambiguously removed using the positional filtering capability of `dplyr`.

Practical Implementation: Removing the Last N Rows (Example 2)

Next, we demonstrate Method 2 by removing the last four rows from the original eight-row data frame. This is useful for truncating larger datasets where a defined number of trailing observations need to be systematically eliminated. We will reset the data frame to its original state and then apply the filter with the subtraction factor set to 4.

In this scenario, `n()` remains 8 (initially). The filtering condition `row_number() <= 8 - 4` simplifies to `row_number() <= 4`. This means only rows 1, 2, 3, and 4 will be retained, while rows 5, 6, 7, and 8 will be discarded. This operation efficiently handles the removal of a block of trailing data.

library(dplyr)

```
#remove last four rows from data frame (assuming df is reset to 8 rows)
```

```
df <- df %>% filter(row_number() <= n()-4)
```

```
#view updated data frame
```

```
df
```

```
team points assists
```

```
1 A 18 5
2 A 13 7
3 A 19 17
4 B 14 9
```

The resulting output clearly demonstrates that the data frame has been successfully reduced to only four observations. The rows corresponding to the latter half of the dataset (rows 5 through 8) have been cleanly removed, confirming the reliability and accuracy of the `row_number() <= n()`

- N methodology.

Alternative Techniques for Subset Selection

While using `filter()` with `row_number()` and `n()` is the recommended `dplyr` approach for clarity, it is worth noting other methods available in R for positional subsetting, particularly for those transitioning from base R. Base R allows subsetting using negative indices to drop specific rows. For example, to remove the last row, one could determine the number of rows using `nrow(df)` and then use a negative index: `df`.

Another powerful `dplyr` function that performs similar positional selection is `slice()`. The `slice()` function is specifically designed to select, retain, or discard rows by integer position. To remove the last row, you could use `df %>% slice(1:(n()-1))`. This syntax achieves the exact same result as the `filter()` method but is arguably slightly more explicit about dealing with row indices, making it a viable alternative for many R programmers.

When choosing between `filter()` and `slice()`, the decision often comes down to personal preference or the complexity of the operation. If the goal is purely positional selection, `slice()` may be preferred due to its dedicated purpose. However, `filter()` is often used globally for conditional logic, and its combination with `row_number()` makes it highly versatile for combining positional and conditional filtering seamlessly.

Conclusion and Best Practices

The use of the `dplyr` package, specifically the combination of `filter()`, `row_number()`, and `n()`, provides the cleanest and most robust method for removing the last row or the last set of N rows from a data frame in R. This approach abstracts away the need for manual row counting and ensures that the code remains highly readable and maintainable.

When implementing these methods in production code, always ensure that your data frame is in the correct order before attempting positional removal. If the data frame is unsorted, removing the "last" row might not correspond to the chronologically or logically last observation. Use functions like `arrange()` before filtering to guarantee that the observations you intend to discard are physically located at the end of the dataset.

Mastering these positional subsetting techniques is crucial for effective data preparation and cleaning in R. By standardizing on `dplyr` methods, developers and analysts can ensure their data pipelines are both efficient and easy to understand, contributing to better collaborative data science practices.

The following tutorials explain how to perform other common functions in `dplyr`: