

Drop Columns by Name in R (With Examples)

Authored by
Mohammed looti

April 4, 2026

RECOMMENDED CITATION

Mohammed looti (2026). *Drop Columns by Name in R (With Examples)*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=3382>

Effective data preparation is paramount in any analytical workflow, and mastering the manipulation of [data frames](#) is a core skill for every user of [R programming](#). As data structures grow in complexity, the necessity to clean, refine, and focus your information becomes critical. One of the most frequently performed data wrangling operations is the removal of superfluous or redundant columns, a technique essential for streamlining your [dataset](#) and improving the efficiency of subsequent analyses.

This comprehensive guide delves into three highly effective and widely adopted methodologies for dropping columns from an [R data frame](#) based on their names. We will examine the foundational approach provided by [Base R](#), the modern, readable syntax offered by the [dplyr package](#), and the high-speed optimization provided by the [data.table package](#).

Understanding these distinct approaches provides tremendous flexibility, allowing analysts to select the method that best aligns with their coding style, performance requirements, and integration within existing data pipelines. By the end of this article, you will be equipped with the knowledge to efficiently manage column structures in any [R data frame](#), regardless of size or complexity.

Three Pillars of Column Management: Base R, `dplyr`, and `data.table`

The [R](#) ecosystem offers multiple robust tools for data manipulation, each with its own advantages. When it comes to removing columns, three methods stand out due to their popularity and efficiency. Choosing between them often depends on the scale of your operation and your familiarity with external libraries.

The traditional method relies on [Base R](#) functionality, specifically the `subset()` function. This approach is universally accessible as it requires no external package installations. The logic is straightforward: columns are excluded by referencing their names within the `select` argument and applying a negative sign. This simplicity makes it an excellent default choice for smaller operations or environments where dependency management is restricted.

#drop col2 and col4 from data frame

```
df_new <- subset(df, select = -c(col2, col4))
```

In contrast, the [dplyr package](#), a fundamental component of the [tidyverse](#) collection, provides a highly expressive and clean grammar for data manipulation. The `select()` function, when combined with the [pipe operator \(%>%\)](#), allows users to chain operations seamlessly. Column removal is achieved using negative indexing on column names within the `select()` call, offering superior readability and making it the preferred method for many modern [R](#) users working in data pipelines.

library(dplyr)

```
#drop col2 and col4 from data frame  
df_new <- df %>% select(-c(col2, col4))
```

Finally, for those dealing with extremely large [datasets](#) where performance and memory efficiency are paramount, the [data.table package](#) is the optimized solution. It enhances standard [data frames](#) into specialized `data.table` objects. Column deletion is distinct in this context, utilizing the `:=` assignment operator to assign [NULL](#) to the unwanted columns. This method typically performs the operation in-place, offering significant speed advantages over methods that create copies of the entire dataset.

library(data.table)

```
#convert data frame to data table  
dt <- setDT(df)  
  
#drop col2 and col4 from data frame  
dt
```

Establishing the Sample Data Frame

To provide clear and reproducible demonstrations of these column dropping techniques, we must first establish a consistent sample [data frame](#). This illustrative dataset, which we will name `df`, mimics typical sports performance statistics, providing several named columns that we can target for removal in the subsequent examples. Using a unified starting point ensures that the results from [Base R](#), [dplyr](#), and [data.table](#) are directly comparable.

The following code block executes the creation of the `df` structure, initializing it with four columns: `team`, `points`, `rebounds`, and `assists`. This creation process utilizes the standard `data.frame()` function available in [Base R](#), confirming our starting data structure before any manipulation begins.

```
#create data frame  
df <- data.frame(team=c('A', 'A', 'B', 'B', 'C', 'C', 'C', 'D'),  
points=c(12, 15, 22, 29, 35, 24, 11, 24),  
rebounds=c(10, 4, 4, 15, 14, 9, 12, 8),  
assists=c(7, 7, 5, 8, 19, 14, 11, 10))  
  
#view data frame  
df
```

```
team points rebounds assists
```

```
1 A 12 10 7
2 A 15 4 7
3 B 22 4 5
4 B 29 15 8
5 C 35 14 19
6 C 24 9 14
7 C 11 12 11
8 D 24 8 10
```

Method 1: Dropping Columns Using Base R's `subset()`

The most accessible method for column removal leverages [Base R](#) capabilities, ensuring that your code is portable and requires no external dependencies. The `subset()` function is specifically designed for straightforward data selection and exclusion. For this practical demonstration, our goal is to eliminate the `points` and `assists` columns, retaining only the identifiers (`team`) and the `rebounds` data in the resulting [data frame](#), which we name `df_new`.

The core concept here is [negative indexing](#) within the `select` argument. By placing a minus sign before the vector of column names (`-c(points, assists)`), we instruct R to exclude those specified columns. This operation is non-destructive; it generates a completely new [data frame](#), leaving the original `df` unaltered, which is a crucial practice for maintaining data integrity and reproducibility in analytical tasks.

#create new data frame by dropping points and assists columns

```
df_new <- subset(df, select = -c(points, assists))
```

```
#view new data frame
```

```
df_new
```

```
team rebounds
```

```
1 A 10
2 A 4
3 B 4
4 B 15
5 C 14
6 C 9
7 C 12
8 D 8
```

As the resulting output clearly verifies, the `df_new` object contains only the `team` and `rebounds` columns. This successfully demonstrates how efficiently [Base R's `subset\(\)`](#) handles column exclusion using a concise, built-in syntax, making it reliable for foundational data cleaning tasks.

Method 2: Dropping Columns with `dplyr`'s `select()`

For users committed to the [tidyverse](#) philosophy, the [`dplyr` package](#) offers a modern, highly consistent, and readable approach to data manipulation. The primary function for column management in this package is `select()`, which seamlessly integrates with the [pipe operator \(`%>%`\)](#) to build fluid data transformation workflows.

To mirror the previous example, we load the [`dplyr` library](#) and then utilize the `select()` function. Just like the Base R method, column deletion is signaled by placing a negative sign before the column names (`-c(points, assists)`). The use of the pipe operator directs the original `df` into the `select()` function, resulting in a new, streamlined [data frame](#).

The key advantage of the `select()` function is its intuitive syntax, which enhances code maintainability and clarity, particularly when the column dropping operation is just one step in a sequence of many transformations, such as grouping, summarizing, or filtering the data.

`library(dplyr)`

```
#create new data frame by dropping points and assists columns
df_new <- df %>% select(-c(points, assists))
```

```
#view new data frame
df_new
```

```
team rebounds
1 A 10
2 A 4
3 B 4
4 B 15
5 C 14
6 C 9
7 C 12
8 D 8
```

The output confirms the successful removal of the `points` and `assists` columns, yielding the exact same structure as the Base R example. This validates the power of the `dplyr` syntax in achieving common data manipulation goals with enhanced readability.

Method 3: High-Performance Deletion with `data.table`

When working with massive [datasets](#), the time and memory overhead associated with copying large objects can become a bottleneck. The [data.table package](#) addresses this challenge by providing optimized operations that often modify the data structure in-place, thus minimizing memory usage and maximizing execution speed.

The process begins by converting our standard `df` into a `data.table` object using `setDT(df)`. Crucially, column deletion in [data.table](#) is achieved by assigning the special value `NULL` to the targeted columns (`points` and `assists`) via the fast assignment operator, `:=`. This assignment tells `data.table` to remove these columns immediately and efficiently.

Unlike the two previous methods that returned a new object, this operation modifies the `dt` object directly. This in-place modification is a hallmark of the [data.table](#) philosophy, making it the preferred tool for performance-critical large-scale data processing in [R](#).

library(data.table)

```
#convert data frame to data table  
dt <- setDT(df)
```

```
#drop points and assists columns  
dt
```

```
#view updated data table  
dt
```

```
team rebounds
```

```
1: A 10
```

```
2: A 4
```

```
3: B 4
```

```
4: B 15
```

```
5: C 14
```

```
6: C 9
```

```
7: C 12
```

```
8: D 8
```

The resulting `dt` object confirms the swift and successful removal of the target columns. Note the slight difference in row numbering (1: instead of 1), which indicates that the structure is now a `data.table` rather than a standard [data frame](#), showcasing its optimized structure ready for further high-speed operations.

Choosing the Optimal Method for Your Data Workflow

While all three presented methods--[Base R's `subset\(\)`](#), [dplyr's `select\(\)`](#), and `data.table`'s `:=NULL`--achieve the identical result of dropping columns by name, the choice of which to implement should be guided by specific project constraints and personal preferences. Key factors include the size of the [dataset](#), the requirement for external package dependencies, and the overall emphasis on code readability.

If your primary concern is maintaining a minimal dependency footprint and working only with built-in functions, the **Base R** method is perfectly adequate and reliable for standard data operations. Conversely, if you prioritize code that is easy to read, debug, and integrates well into a modern, sequential data processing structure, `dplyr` provides the most intuitive and clean syntax, especially when leveraging the piping mechanism common throughout the [tidyverse](#).

However, when computational efficiency and memory optimization become critical--typically when processing data structures exceeding several million rows--the speed advantages offered by `data.table` are unmatched. Its C-based backend and in-place assignment capabilities make it the definitive choice for big data analysis in the [R](#) environment. Selecting the right tool ensures not only correct results but also an efficient and scalable analytical process.

Further Exploration and Essential Data Manipulation Resources

Dropping columns is merely one facet of comprehensive data wrangling in [R](#). To become truly proficient in data analysis, it is essential to master a range of manipulation techniques that allow you to reshape, filter, and transform data frames according to analytical needs. We strongly encourage you to continue expanding your knowledge of data handling functions within the [tidyverse](#) and **Base R**.

Deepening your expertise in related tasks such as column selection, renaming, and row filtering will significantly enhance your ability to prepare complex data for modeling and visualization. These skills collectively form the foundation for robust and reliable statistical analysis.

Consider these additional tutorials to build upon your knowledge of column manipulation in **R**:

How to [Select Columns in R](#)

How to [Rename Columns in R](#)

How to [Filter Rows in R](#)