

Merge Multiple Data Frames in R (With Examples)

Authored by
Mohammed looti

November 2, 2025

RECOMMENDED CITATION

Mohammed looti (2025). *Merge Multiple Data Frames in R (With Examples)*.
PSYCHOLOGICAL STATISTICS. Retrieved from
<https://statistics.arabpsychology.com/?p=8382>

When working with complex datasets in the [R programming language](#), a common requirement is consolidating information scattered across multiple source files or objects. This necessitates merging several [data frames](#) into a single, cohesive structure. Fortunately, R offers robust and efficient tools for this task, primarily relying on two powerful methodologies: utilizing core **Base R** functions or leveraging the streamlined approach provided by the **Tidyverse** ecosystem.

This tutorial will explore both methods, detailing the required syntax, explaining the underlying mechanics, and providing practical examples to help you master the art of merging data frames in R.

Overview of Merging Strategies in R

Regardless of the approach you choose, the goal remains the same: iteratively combine a list of data frames based on shared key columns. Below is a quick comparison of the foundational code required for each strategy.

Method 1: Utilizing Base R Functions

The Base R approach relies on the standard `merge()` function, but since `merge()` can only handle two objects at a time, we use the higher-order function `Reduce()` to apply `merge()` sequentially across a list of data frames.

1. Collect all data frames into a single list

```
df_list <- list(df1, df2, df3)
```

2. Use Reduce() to apply merge() iteratively

```
Reduce(function(x, y) merge(x, y, all=TRUE), df_list)
```

Method 2: Streamlining with Tidyverse

The [Tidyverse](#) provides a collection of packages optimized for data science. Specifically, the `purrr::reduce()` function combined with the `dplyr::full_join()` function offers a clean, highly readable syntax for achieving the same result.

library(tidyverse)

1. Collect all data frames into a list

```
df_list <- list(df1, df2, df3)
```

2. Use the pipe operator and reduce with full_join

```
df_list %>% reduce(full_join, by='variable_name')
```

The subsequent sections detail how to implement each of these methods using realistic data examples, ensuring a complete understanding of the underlying logic.

Method 1: Leveraging Base R Functions (The Reduce Approach)

The standard `merge()` function in R is designed for binary operations--joining two data frames. When faced with three or more data frames, we must find a way to apply this function repeatedly. This is where the powerful `Reduce()` function comes into play. The `Reduce()` function takes a binary function and applies it sequentially to the elements of a list, accumulating the result.

In the context of merging, we pass `merge()` as the binary function, ensuring that the output of the first merge (`df1 + df2`) becomes the first argument for the next merge (`result + df3`), and so on. The argument `all=TRUE` ensures that we perform a **full outer join**, retaining all rows present in any of the original data frames.

Detailed Example: Merging Data Frames Using Base R

For this example, suppose we are tracking financial data where revenue, expenses, and profit are recorded in separate data frames, keyed by a common `id` variable. Notice that not all IDs are present in every data frame, which makes the full outer join crucial.

We first define our sample data frames:

```
# Define data frame 1: Revenue data
```

```
df1 <- data.frame(id=c(1, 2, 3, 4, 5),  
revenue=c(34, 36, 40, 49, 43))
```

```
# Define data frame 2: Expense data
```

```
df2 <- data.frame(id=c(1, 2, 5, 6, 7),  
expenses=c(22, 26, 31, 40, 20))
```

```
# Define data frame 3: Profit data
```

```
df3 <- data.frame(id=c(1, 2, 4, 5, 7),  
profit=c(12, 10, 14, 12, 9))
```

To consolidate this information, we first place them into a list and then apply the `Reduce()` function. This ensures that the data frames are combined one after the other, respecting the full set of IDs across all sources.

```
# 1. Combine all data frames into a list object
```

```
df_list <- list(df1, df2, df3)
```

```
# 2. Execute the iterative merge using Reduce()
Reduce(function(x, y) merge(x, y, all=TRUE), df_list)
```

```
id revenue expenses profit
1 1 34 22 12
2 2 36 26 10
3 3 40 NA NA
4 4 49 NA 14
5 5 43 31 12
6 6 NA 40 NA
7 7 NA 20 9
```

The resulting data frame successfully includes all unique `id` values (1 through 7). Where a specific ID lacked data in one of the original frames (e.g., ID 3 had no expense or profit data), the resulting merged column is correctly populated with **NA** (Not Available) values, which is the standard behavior for a full outer join.

Method 2: Streamlining Merges with the Tidyverse

For those accustomed to the Tidyverse suite--which emphasizes clear, pipeline-based workflows--merging multiple data frames is often simpler and more intuitive. This method leverages two key functions from Tidyverse packages: `purrr::reduce()` and `dplyr::full_join()`.

The `reduce()` function from the `purrr` package operates identically to the Base R `Reduce()`, applying a function sequentially. However, instead of using the generic `merge()`, we use `full_join()` from the `dplyr` package. The advantage here is clarity; `full_join()` explicitly states the type of join being performed, and Tidyverse join functions are typically more performant on large datasets due to their optimized C++ backend.

Detailed Example: Merging Data Frames Using Tidyverse

We will use the exact same input data frames (`df1`, `df2`, and `df3`) to demonstrate that the Tidyverse method yields an identical, consistent result, while offering improved syntax and potential speed benefits.

First, we ensure the data frames are defined as before:

```
# Define data frames identical to Method 1
df1 <- data.frame(id=c(1, 2, 3, 4, 5),
revenue=c(34, 36, 40, 49, 43))
```

```
df2 <- data.frame(id=c(1, 2, 5, 6, 7),
  expenses=c(22, 26, 31, 40, 20))
```

```
df3 <- data.frame(id=c(1, 2, 4, 5, 7),
  profit=c(12, 10, 14, 12, 9))
```

The Tidyverse approach requires loading the suite and then piping the list of data frames directly into `reduce()`, specifying `full_join` as the function to be applied and `'id'` as the joining key.

library(tidyverse)

```
# 1. Collect all data frames into a list
```

```
df_list <- list(df1, df2, df3)
```

```
# 2. Pipe the list into reduce(), using full_join and specifying the 'id' key
```

```
df_list %>% reduce(full_join, by='id')
```

```
id revenue expenses profit
1 1 34 22 12
2 2 36 26 10
3 3 40 NA NA
4 4 49 NA 14
5 5 43 31 12
6 6 NA 40 NA
7 7 NA 20 9
```

As demonstrated by the output, the result is identical to the Base R method, achieving a complete consolidation of data across all unique identifiers.

Performance and Best Practices

While both Base R and Tidyverse methods achieve the desired result, data scientists often face the question of which approach is preferable. The choice typically depends on the size of the data and the user's familiarity with the respective package ecosystems.

Base R (`Reduce(merge)`): This method is essential for environments where external packages cannot be installed or used. It is robust and relies only on the core R installation. However, the performance can degrade more rapidly than Tidyverse methods when dealing with extremely large data frames (millions of rows).

Tidyverse (`reduce(full_join)`): This approach is generally recommended for modern data

workflows. The code is more readable, and the underlying functions in `dplyr` are highly optimized for speed and memory efficiency, making this method noticeably quicker when working with **extremely large data frames**.

A key best practice when merging multiple data frames is to consistently name the key variables (e.g., always use `id`, not `customer_id` in one frame and `client_id` in another). This eliminates ambiguity and ensures that both the Base R and Tidyverse methods can correctly identify the columns upon which to join.

Conclusion and Further Learning

Merging multiple data frames is a fundamental data manipulation skill in R. Whether you opt for the iterative power of `Reduce` combined with Base R's `merge()`, or the speed and clarity offered by the Tidyverse combination of `reduce()` and `full_join()`, R provides effective solutions for consolidating complex data structures.

Mastering these techniques will significantly enhance your ability to clean, prepare, and analyze data efficiently. For further development of your R data manipulation skills, consider exploring specialized documentation on different join types (e.g., inner, left, semi-joins) and advanced data restructuring functions.

Additional Resources

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