

How to Easily Create Grouped Frequency Tables in R

Authored by
stats writer

December 4, 2025

RECOMMENDED CITATION

stats writer (2025). *How to Easily Create Grouped Frequency Tables in R*.
PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=104637>

Generating a frequency table by group is a fundamental requirement in data analysis, allowing statisticians and analysts to quickly ascertain the distribution of observations across specific categories or subsets of the data. While base R provides functions like `table()` and `aggregate()` for this purpose, modern data wrangling in R often relies on the highly efficient and readable methods provided by the Tidyverse ecosystem, particularly the `dplyr` package.

This tutorial focuses on leveraging the powerful combination of `group_by()` and `summarize()` from `dplyr`. This approach streamlines the process of calculating frequency counts for combinations of categorical variables, producing clean, organized results that are easy to interpret and further manipulate. By structuring the analysis around these functions, we move beyond manual manipulation toward reproducible and scalable data pipelines.

Understanding Frequency Tables and Grouping in Data Analysis

A frequency table provides a succinct summary of the occurrences of distinct values within a dataset. When we introduce grouping, we are asking R to calculate these frequencies not just for the entire dataset, but separately for observations that share common characteristics defined by one or more grouping variables. This is crucial for segmentation and comparative analysis--for example, analyzing product popularity segregated by region, or, as in our example, examining player positions segregated by team.

Effective grouping transforms raw data into actionable insights. In essence, the process involves three key conceptual steps: defining the groups, performing the calculation (counting observations), and presenting the results in a new, summarized structure. The efficiency and syntax of the `dplyr` package make executing these steps straightforward, establishing it as the preferred tool for such operations within the R environment.

The Role of the dplyr Package in R Data Wrangling

The `dplyr` package is the cornerstone of data manipulation in R's Tidyverse. It offers a standardized set of verbs designed to make data handling intuitive and expressive. When dealing with grouped calculations, `dplyr` excels because its functions are optimized for vectorized operations, ensuring performance even with large datasets. Furthermore, the use of the pipe operator (`%>%`) allows users to chain operations together logically, significantly enhancing the readability of the code.

Using `dplyr` minimizes the need for complex nested functions or iterative loops, which are often difficult to debug and maintain. Instead, analysts can write code that clearly reflects the analysis steps: first, define the data source, then define the grouping, and finally, define the summary calculation. This structured approach is highly valued in collaborative and professional data science environments.

Essential dplyr Functions for Grouped Summarization

To achieve grouped frequency counts, we rely on three primary components within the `dplyr` framework. The first component is the `group_by()` function. This function takes an existing data structure and flags specific columns as grouping variables. All subsequent operations applied to this grouped object will then be performed separately within the scope of those defined groups.

The second essential component is the `summarize()` function (or `summarise()`, which is an alias). This function collapses the grouped data into a single row for each unique group combination, calculating specified metrics. In the context of frequency tables, the metric we need is a count of observations within each group. This leads us to the final component: the special helper function, `n()`, which is used exclusively within `summarize()` to count the number of rows (observations) in the current group.

By combining these tools--first applying `group_by()` to set the scope, and then using `summarize()` with `n()` to execute the count--we efficiently generate the desired grouped frequency table. The following R code snippet illustrates this standard syntax:

You can use the following functions from the `dplyr` package to create a frequency table by group in R:

library(dplyr)

```
df %>%  
group_by(var1, var2) %>%  
summarize(Freq=n())
```

The following practical example shows how to apply this syntax to a sample dataset.

Setting Up the R Environment and Initial Data Frame

Before executing any data manipulation using `dplyr`, it is necessary to ensure the package is loaded into the R session using the `library()` function. Once loaded, we can proceed to define our sample data. For this demonstration, we will create a simple data frame that simulates a roster of players, detailing their team assignment and position. This data frame serves as the input for our frequency analysis.

The structure of the data frame, containing categorical variables like `team` and `position`, is ideal for demonstrating grouped frequency counts. We are interested in finding out how many players occupy each position, but segmented by their respective teams (A and B). This requires grouping by both variables simultaneously.

```
#create data frame
df <- data.frame(team=c('A', 'A', 'A', 'A', 'B', 'B', 'B', 'B'),
position=c('G', 'G', 'G', 'F', 'G', 'F', 'F', 'C'))
```

```
#view data frame
```

```
df
```

```
team position
```

```
1 A G
```

```
2 A G
```

```
3 A G
```

```
4 A F
```

```
5 B G
```

```
6 B F
```

```
7 B F
```

```
8 B C
```

Our objective is to create a frequency table that shows the count of each unique position, categorized by the corresponding team.

Example: Creating the Grouped Frequency Table

To execute the grouped frequency analysis, we initiate the pipeline by feeding the data frame `df` into the pipe operator. Next, we call `group_by()`, specifying both `team` and `position` as the grouping variables. This instructs R to treat every unique combination of team and position (e.g., Team A, Position G; Team B, Position F) as a distinct unit for calculation. Finally, we apply `summarize()`, creating a new column named `Freq` which holds the count of observations (rows) in each group, calculated using `n()`.

This single, cohesive command structure is highly efficient. It avoids the manual sorting and counting that older methods might require. The output produced by `summarize()` will be a new tibble (a modern data frame structure) that lists every unique group combination and its corresponding frequency count.

library(dplyr)

```
#calculate frequency of position, grouped by team
```

```
df %>%
```

```
group_by(team, position) %>%
```

```
summarize(Freq=n())
```

```
# A tibble: 5 x 3
# Groups: team
team position Freq

1 A F 1
2 A G 3
3 B C 1
4 B F 2
5 B G 1
```

Interpreting the Grouped Frequency Output

The result is a tibble showing five unique combinations of team and position, along with the calculated frequency (count). Note that there are only five rows, even though the original data frame had eight, because the `summarize()` function aggregated identical rows. Analyzing this output allows us to derive specific insights about the composition of each team. The `Groups: team` line indicates that the resulting tibble is still grouped by `team`, which is important if further grouped operations were planned.

The interpretation of each row is straightforward, quantifying exactly how many players in a given team hold a specific position. This clear and concise representation is the primary benefit of using grouped summarization for frequency analysis.

Here's how to interpret the output:

- 1 player on team A has position 'F' (Forward)
- 3 players on team A have position 'G' (Guard)
- 1 player on team B has position 'C' (Center)
- 2 players on team B have position 'F' (Forward)
- 1 player on team B has position 'G' (Guard)

Customizing the Output: Renaming the Frequency Column

It is important to use descriptive column names in any analytical output to ensure clarity for others viewing the results. Although `Freq` is commonly understood, renaming the count column can improve overall project readability, especially when combining multiple summary statistics. We rename the column simply by changing the variable name assigned within the `summarize()` function call.

For example, if the frequency represents a count of total individuals, using `count` or `Total_Players` might be more expressive than `Freq`. This small adjustment is performed directly

within the argument list of `summarize()`, demonstrating the flexibility and ease of customization provided by `dplyr`.

For example, we could rename the column to be named 'count' instead:

library(dplyr)

```
#calculate frequency of position, grouped by team
df %>%
group_by(team, position) %>%
summarize(count=n())
```

```
# A tibble: 5 x 3
# Groups: team
team position count
```

```
1 A F 1
2 A G 3
3 B C 1
4 B F 2
5 B G 1
```

Further Applications and Conclusion

Generating grouped frequency tables is often just the initial step in a larger analytical workflow. Once the counts are calculated, analysts might proceed to convert these counts into relative frequencies (percentages) by further manipulating the summarized table. This involves applying window functions or additional `mutate()` operations, typically grouped by the higher-level variable (e.g., `team`) to calculate the percentage contribution of each position within its respective team total.

The methods demonstrated here, utilizing `group_by()` and `summarize()`, represent the most robust and standard approach for conducting segmented frequency analysis in R. Mastering these core `dplyr` concepts ensures that your data wrangling processes are clean, efficient, and easily scalable for handling complex, large-scale datasets.

Related Data Wrangling Tutorials

The following tutorials explain how to perform other common data manipulation functions using `dplyr`, building upon the foundational knowledge presented here: