

How to Easily Remove NAs from ggplot2 Plots

Authored by
stats writer

November 22, 2025

RECOMMENDED CITATION

stats writer (2025). *How to Easily Remove NAs from ggplot2 Plots*. PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=99844>

The handling of NA values, or missing data, is a critical step in any data analysis and visualization workflow, particularly when utilizing powerful statistical graphics libraries like `ggplot2` in R. While `ggplot2` often attempts to accommodate missing observations gracefully, these null entries can sometimes clutter visualizations or lead to misleading interpretations, especially in categorical plots like bar charts. The primary goal of this guide is to demonstrate the robust, recommended method for explicitly excluding these unwanted NA values from your plots by preprocessing the underlying data structure before rendering the graphic.

In `ggplot2`, you might encounter the internal `na.rm` parameter within certain geometric functions (e.g., `geom_point(na.rm=TRUE)`). While this parameter is useful for removing missing coordinates during the plotting process (e.g., removing a point if either X or Y is missing), it typically does not address the categorical representation of missing data itself, such as an explicit '<NA>' category appearing on an axis label in a bar plot. Therefore, to ensure a truly clean visualization free of NA values, we must manipulate the source data frame using logical indexing, which offers greater control and clarity over the final output.

Introduction to Missing Data Handling in ggplot2

Missing data is an inevitable challenge in real-world datasets, and how analysts choose to handle it profoundly impacts the resulting visualizations. When using `ggplot2`, the default behavior for categorical variables containing NA values is often to treat 'NA' as a valid level or category, resulting in a specific bar or grouping dedicated solely to these null entries. While this behavior is technically correct--it alerts the user to the presence of missing data--it is frequently undesirable when the goal is to focus exclusively on the distribution of known, observed categories. Analysts often seek a method to seamlessly exclude these rows entirely before the plotting function is called, ensuring the visual focus remains strictly on complete observations.

The standard R package approach involves leveraging powerful base functions designed for data manipulation. Specifically, we combine the logical inverse operator (!) with the specialized `is.na()` function within the context of the `subset()` function. This technique allows us to filter the input data frame dynamically, providing the plotting function with a sanitized dataset that contains only non-missing values in the target column. This approach is highly flexible and transparent, making it the preferred method for maintaining data integrity while producing publication-quality graphics.

To implement this, the syntax shown below provides a general template. It demonstrates how to utilize the `subset` command directly within the `ggplot()` call, effectively streamlining the data preparation and visualization steps. This pattern ensures that any temporary filtering is applied only for the visualization context, leaving the original source data frame intact and unchanged for subsequent analyses.

You can use the following basic syntax to remove NA values from a plot in `ggplot2`:

library(ggplot2)

```
ggplot(data=subset(df, !is.na(this_column)), aes(x=this_column)) +  
geom_bar()
```

This particular example creates a bar plot and removes any rows in the data frame where an NA value occurs in the column called **this_column**. The core mechanism is the expression `!is.na(this_column)`, which returns a logical vector of **TRUE** for non-missing entries and **FALSE** for missing entries. The subset() function then uses this vector to retain only those rows where the condition is **TRUE**.

Understanding na.rm vs. Data Subsetting

When newcomers to ggplot2 encounter missing data, they often first look for the **na.rm** argument, following the convention seen in other R functions like `mean()` or `sum()`. It is crucial to understand that within ggplot2, the `na.rm = TRUE` argument is typically intended for specific geometric layers (like `geom_point()`, `geom_line()`, or `geom_smooth()`) and primarily governs how the geometry handles missing values in continuous coordinate mapping. If a point is defined by `x=NA` or `y=NA`, `na.rm=TRUE` tells the layer to silently drop that specific observation from the geometric drawing phase. This is effective for continuous data representations but often fails to remove the explicit categorical '<NA>' representation on the axis itself, which is often derived from the factor levels of the underlying data.

Data Subsetting, using functions like the base R subset() function or the Tidyverse `filter()` function, operates at a fundamental level: the data preparation stage. By filtering the data frame itself, we ensure that the ggplot2 mapping functions never even encounter the missing rows for the specific column of interest. This proactive removal is significantly cleaner and more reliable when dealing with categorical axes. It guarantees that the resulting plot reflects only the data that passed the quality check, thus eliminating the '<NA>' bar or category label entirely, preventing confusion and enhancing the visual clarity of the analysis.

Therefore, while `na.rm` serves a specific purpose in smoothing out plotting artifacts in continuous geometries, the gold standard for removing categorical NA values from axes or legends is always pre-processing the data. The subsequent example will highlight this difference by showing how a bar plot, which counts occurrences based on factors, inherently includes the NA category unless the data is explicitly filtered using the `!is.na()` condition.

The Standard Approach: Using the subset() and !is.na() Functions

The combination of the subset() function and the is.na() function represents the most common and

robust way to manage missing observations in base R data processing before visualization. The logic is straightforward yet powerful: first, we identify all elements in the target column that are missing using `is.na(column)`. This returns a logical vector (TRUE/FALSE) where TRUE indicates a missing value. Second, we apply the logical negation operator (!) to this vector, flipping all TRUES to FALSEs and vice-versa. This new vector now indicates **TRUE** only for non-missing, valid entries.

The resulting logical vector is then passed as the condition argument to the `subset()` function. The `subset()` function performs the selection, retaining only those rows in the data frame for which the logical condition evaluated to TRUE. By executing this operation directly within the `ggplot()` call, we create a temporary, clean dataset that feeds directly into the aesthetic mapping (`aes()`), thereby ensuring the plot is generated using only the desired, complete data points. This methodological precision is key to generating accurate and unambiguous graphical summaries.

While the example uses base R functions for maximum compatibility, those working within the Tidyverse ecosystem might prefer the equivalent syntax using the `dplyr::filter()` function, which achieves the same filtering effect: `ggplot(data = df %>% filter(!is.na(team)), aes(x=team)) + geom_bar()`. Regardless of the specific functions used, the underlying principle remains constant: the data must be scrubbed of missing entries in the relevant aesthetic column before the visualization engine processes it. This meticulous attention to data quality ensures that the visualizations accurately reflect the population of observed values rather than confusing the interpretation with unknown or null entries.

Detailed Walkthrough: Setting Up the Example Dataset

To illustrate this process clearly, let us establish a simple scenario involving sports statistics where some data entries are incomplete. We will create a sample data frame named **df** that tracks the points scored by players, but crucially, includes several rows where the corresponding team identity is missing, represented by **NA**. This mimics common data collection issues where categorical identifiers are occasionally absent.

This dataset contains two key variables: `team`, which is the categorical variable of interest that contains missing values, and `points`, which is a quantitative measure. The objective is to visualize the frequency count of observations per team. If we simply pass the unfiltered data to `ggplot2`, we will see how it naturally incorporates the two missing team entries into the resulting bar chart, which we aim to prevent. The following code demonstrates the creation and structure of this illustrative data frame:

The following example shows how to use this syntax in practice. Suppose we have the following data frame that contains information on the number of points scored by basketball players on various teams:

```
#create data frame
df <- data.frame(team=c('A', 'A', NA, NA, 'B', 'B', 'B', 'B'),
points=c(22, 29, 14, 8, 5, 12, 26, 36))
```

```
#view data frame
```

```
df
```

```
team points
```

```
1 A 22
```

```
2 A 29
```

```
3 <NA> 14
```

```
4 <NA> 8
```

```
5 B 5
```

```
6 B 12
```

```
7 B 26
```

```
8 B 36
```

As observed in the output, rows 3 and 4 explicitly show <NA> under the `team` column. When a bar plot is generated, the `geom_bar()` layer, by default, calculates the count of each unique value in the aesthetic mapping (`aes(x=team)`). Since '<NA>' is a unique, albeit null, observation, it will be counted alongside 'A' and 'B', leading to an extra bar in our visualization. We proceed to plot this raw data to clearly establish the problem we are attempting to solve with data subsetting.

Visualizing Raw Data: The Impact of Missing Values on Bar Charts


We now attempt to generate the initial bar plot using the unfiltered `df`. This visualization is intended to show the frequency distribution of observations across the existing teams. Because we are using `geom_bar()` without any aggregation, `ggplot2` automatically performs a count of the rows for each unique factor level specified in the x-aesthetic. Since the `team` column contains two instances of missing data, the resulting plot will allocate a specific visual element (a bar) to represent these null entries, treating them as a distinct category.

This default inclusion of NA values as a category is often problematic for reporting, as it forces the reader to acknowledge data that is inherently unclassified or incomplete, distracting from the main analysis of observed teams 'A' and 'B'. Furthermore, in certain contexts, visualizing the NA count separately may unnecessarily compress the scale for the valid categories, diminishing the clarity of the primary data trends.

Now suppose we attempt to create a bar plot in `ggplot2` to visualize the number of occurrences of each team:

library(ggplot2)

```
#create bar plot to visualize occurrences by team  
ggplot(df, aes(x=team)) +  
geom_bar()
```



Notice that the plot automatically creates a bar to display the occurrences of NA values in the **team** column. The height of this bar corresponds exactly to the two rows identified as missing in our initial data frame setup. The next step involves applying our filtering strategy using the subset() function to remove these rows before the plot is rendered.

Implementing the Solution: Removing NAs Using Data Subsetting

To achieve a cleaner plot that only displays the counts for teams 'A' and 'B', we must introduce the filtering condition directly into the `ggplot()` function call. We utilize the subset() function, passing it the entire **df** and applying the condition `!is.na(team)`. This operation generates a temporary,


filtered version of the data frame that completely excludes the rows where the `team` column contains an NA. This filtered data frame is then passed as the `data` argument to the `ggplot()` function.

The resulting plot, generated from this subsetted data, will now only contain information for valid team entries. The aesthetic mapping (`aes(x=team)`) will only see two categories ('A' and 'B'), and consequently, the visualization will only contain two bars corresponding to their respective counts. This is the most effective and universally applicable method for removing missing categorical values from ggplot2 visualizations, ensuring that the plot focuses entirely on the observed data distribution.

To remove this bar from the plot, we can use the subset() function to subset the data frame to only include rows where the value in the **team** column is not NA:

```
library(ggplot2)
```

```
#create bar plot to visualize occurrences by team and remove NA  
ggplot(data=subset(df, !is.na(team)), aes(x=team)) +  
geom_bar()
```



This bar plot still displays the number of occurrences for the values 'A' and 'B' in the **team** column but it no longer includes a bar to display the number of occurrences for NA values. The counts (2 for A, 4 for B) remain accurate relative to the filtered data set, but the visual field is free from the ambiguity of missing observations.

Alternative Methods for Handling Missing Data in R

While the `subset` approach is excellent for quick, localized filtering within a `ggplot` call, R and its surrounding package ecosystem offer several other powerful methods for handling NA values, especially when cleaning or analysis requires more complex imputation or permanent removal. For instance, the base R function `na.omit(df)` provides a blunt but effective way to remove any row that contains at least one NA across any column. While simple, this method can lead to significant data loss if missingness is common across many different variables, as it removes complete rows based on missingness in unrelated columns.

For more selective or programmatic data cleaning, the Tidyverse packages offer refined tools. Using `dplyr::drop_na(df, team)` explicitly removes only those rows where NA is present in the specified `team` column, offering better control than `na.omit`. Furthermore, for projects requiring

data imputation--estimating missing values rather than removing them--packages like `mice` or methods provided by `tidymodels` are essential. Choosing the right method (removal, imputation, or explicit subsetting for visualization) depends entirely on the analytical goals and the nature of the missing data itself.

Conclusion and Further Reading

Mastering the technique of data subsetting using `!is.na()` is fundamental for any R user creating professional visualizations with `ggplot2`. It moves beyond simple aesthetic adjustments and engages directly with data integrity, ensuring that the visual output is a true reflection of the complete observations available. This methodology ensures consistency and reliability, particularly when dealing with large datasets where missingness is common but needs to be excluded from primary visual summaries.

The practice demonstrated here--filtering the data source before passing it to the plotting function--is transferable across many other R visualization tasks, reinforcing the importance of data pre-processing as a foundational step. By implementing robust filtering techniques, users gain full control over how missing information is handled, leading to more impactful and less ambiguous graphical representations.

The following tutorials explain how to perform other common tasks in `ggplot2`: