# Learning to Visualize Data: A Step-by-Step Guide to Creating Heatmaps in R with ggplot2

Authored by

**Mohammed looti**

November 9, 2025

Data visualization is a critical component of modern data analysis, allowing researchers and analysts to quickly identify patterns and correlations within complex datasets. Among the most powerful tools available for visualizing multivariate data is the **heatmap**. A heatmap represents the magnitude of a phenomenon as color in two dimensions, making it exceptionally effective for displaying correlation matrices or showing the intensity of variables across different observations.

This comprehensive tutorial provides a step-by-step guide on how to construct a high-quality, customizable heatmap in **R**, leveraging the powerful and flexible features of the **ggplot2** package. While the process may seem straightforward, preparing the data correctly is essential, as heatmaps typically require data to be in a specific structure--the long format. We will begin by introducing the dataset and then move through the necessary data preparation steps before diving into the visualization code.

## Prerequisite Steps: Data Preparation and Initial Inspection

To illustrate the process of creating a heatmap, we will utilize the renowned built-in **R** dataset, **mtcars**. This dataset contains 32 observations (cars) and 11 variables related to car performance, offering a rich source of numerical data suitable for heatmap visualization. Before we can use this data for visualization, it is crucial to understand its current structure.

The **mtcars** dataset is currently in a **wide format**, where each variable occupies its own column. The rows represent the car models. We must inspect the initial structure to confirm its layout and the types of variables involved. The following code snippet allows us to view the first few entries of the dataset:

**#view first six rows of *mtcars***
**head(mtcars)**

# mpg cyl disp hp drat wt qsec vs am gear carb
#Mazda RX4 21.0 6 160 110 3.90 2.620 16.46 0 1 4 4
#Mazda RX4 Wag 21.0 6 160 110 3.90 2.875 17.02 0 1 4 4
#Datsun 710 22.8 4 108 93 3.85 2.320 18.61 1 1 4 1
#Hornet 4 Drive 21.4 6 258 110 3.08 3.215 19.44 1 0 3 1
#Hornet Sportabout 18.7 8 360 175 3.15 3.440 17.02 0 0 3 2
#Valiant 18.1 6 225 105 2.76 3.460 20.22 1 0 3 1

For **ggplot2** to correctly map variables onto the axes (car names vs. variable types) and use a third variable for color intensity (the actual value), we must convert this structure into a **long format**. This transformation is fundamental for creating effective statistical graphics using the grammar of graphics approach, particularly for visualizations like heatmaps where multiple

columns need to be stacked into key-value pairs.

## Transforming Data from Wide to Long Format

The process of converting the data structure from **wide format** to **long format** is often referred to as 'melting' the data. This requires the use of specialized tools. We will rely on the **reshape2** package, which provides the highly efficient `melt()` function specifically designed for this purpose. The long format will restructure the data such that we have one column identifying the car (our row label), one column identifying the variable name (our column label), and one column containing the corresponding numerical value (our color intensity).

After loading the necessary library, we apply the `melt()` function to the **mtcars** data frame. Since the car names are stored as row names rather than a standard column, we must explicitly extract these names and assign them to a new column named `car`. This ensures that the car model is preserved as an identifiable observation in the melted data frame, which is essential for labeling the heatmap's y-axis. The process is detailed below:

**#load *reshape2* package to use melt() function**
**library(reshape2)**

#melt mtcars into long format
melt_mtcars <- melt(mtcars)

#add column for car name
melt_mtcars$car <- rep(row.names(mtcars), 11)

#view first six rows of *melt_mtcars*
head(melt_mtcars)

# variable value car
#1 mpg 21.0 Mazda RX4
#2 mpg 21.0 Mazda RX4 Wag
#3 mpg 22.8 Datsun 710
#4 mpg 21.4 Hornet 4 Drive
#5 mpg 18.7 Hornet Sportabout
#6 mpg 18.1 Valiant

The resulting data frame, `melt_mtcars`, is now properly formatted for visualization. The `variable` column will define the heatmap's x-axis, the `car` column will define the y-axis, and the `value` column will determine the color intensity of each cell, fulfilling the requirements for creating our initial **heatmap** using **ggplot2**'s geometry layers.

## Initial Heatmap Visualization and Scale Challenges

With the data successfully transformed, we can proceed to generate the first iteration of the heatmap. The core of the **ggplot2** implementation relies on mapping the aesthetics: `variable` to the x-axis, `car` to the y-axis, and `value` to the fill color. We use `geom_tile()` to draw the rectangular tiles that form the heatmap, and `scale_fill_gradient()` to define the continuous color scale, ranging from white (low values) to red (high values).

The following code block executes the initial visualization:

**library(ggplot2)**

ggplot(melt_mtcars, aes(variable, car)) +
geom_tile(aes(fill = value), **colour = "white") +**
**scale_fill_gradient(low = "white", high = "red")**

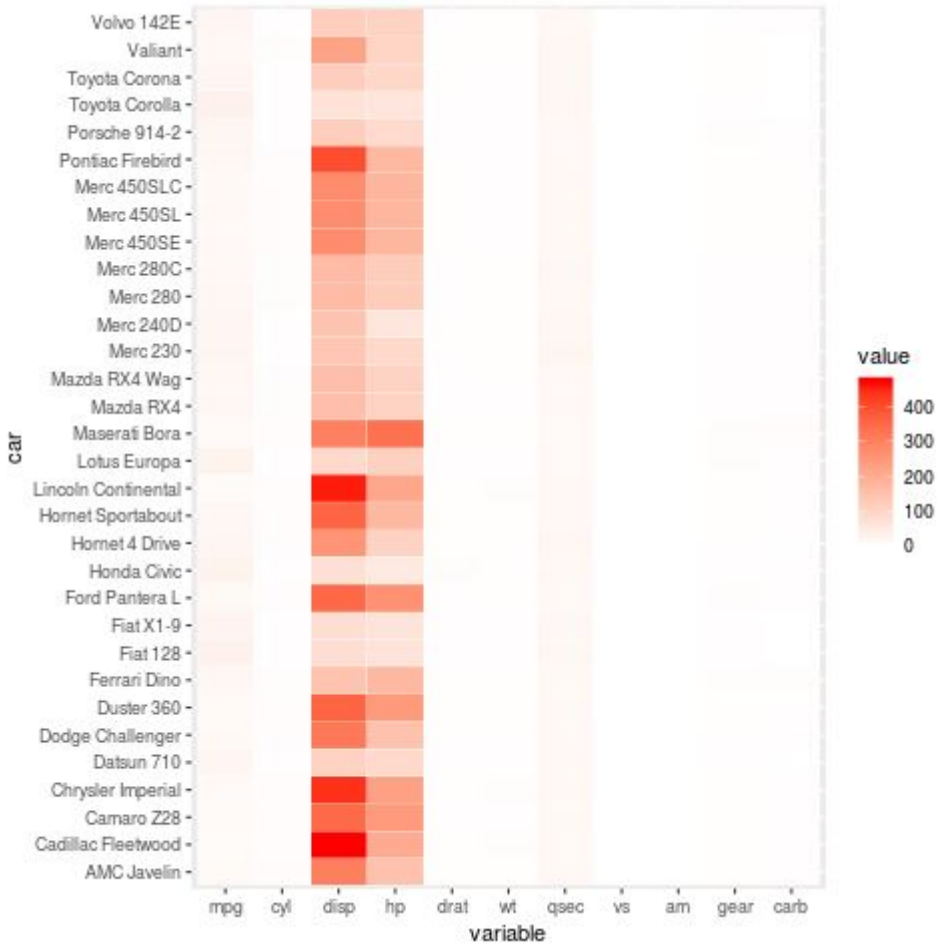


Upon reviewing the initial output, a significant issue related to scale disparity becomes immediately

apparent. Variables such as *disp* (displacement) have numerical values much larger than those of variables like *mpg* or *drat*. Because the color scale is applied across the entire dataset, these high-magnitude variables dominate the visualization, causing most other cells in the heatmap to appear uniformly light (white), making the subtle variations within the smaller-valued variables invisible. This scaling problem severely limits the analytical utility of the visualization, as the color gradient is effectively only showing variation in the largest numerical column.

## Resolving Scale Disparity through Data Rescaling

To create a truly informative **heatmap**, we must normalize the data so that the color intensity reflects the relative position of a value within its specific variable column, rather than its absolute magnitude compared to all other variables. This technique, known as **data rescaling** or normalization, typically involves transforming all values within a column to fit within a standardized range, commonly 0 to 1.

We will achieve this using two specific **R** packages: **scales**, which provides the necessary `rescale()` function, and **plyr**, which offers the powerful `ddply()` function for applying transformations to subsets of a data frame (in this case, applying the rescaling function column-by-column). By using `ddply()` grouped by the `variable` column, we ensure that the rescaling operates independently on each measured characteristic (e.g., all *mpg* values are scaled together, and all *hp* values are scaled together).

The implementation involves loading these libraries and then using `ddply()` to create a new column, `rescale`, based on the normalized `value` column. This transformed data frame is then used to generate a much more insightful heatmap, as shown in the code below:

**#load libraries**
**library(plyr)**
**library(scales)**

#rescale values for all variables in melted data frame
melt_mtcars <- ddply(melt_mtcars, .(variable), transform, rescale = rescale(value))

#create heatmap using rescaled values
ggplot(melt_mtcars, aes(variable, car)) +
geom_tile(aes(fill = rescale), colour = "white") +
scale_fill_gradient(low = "white", high = "red")

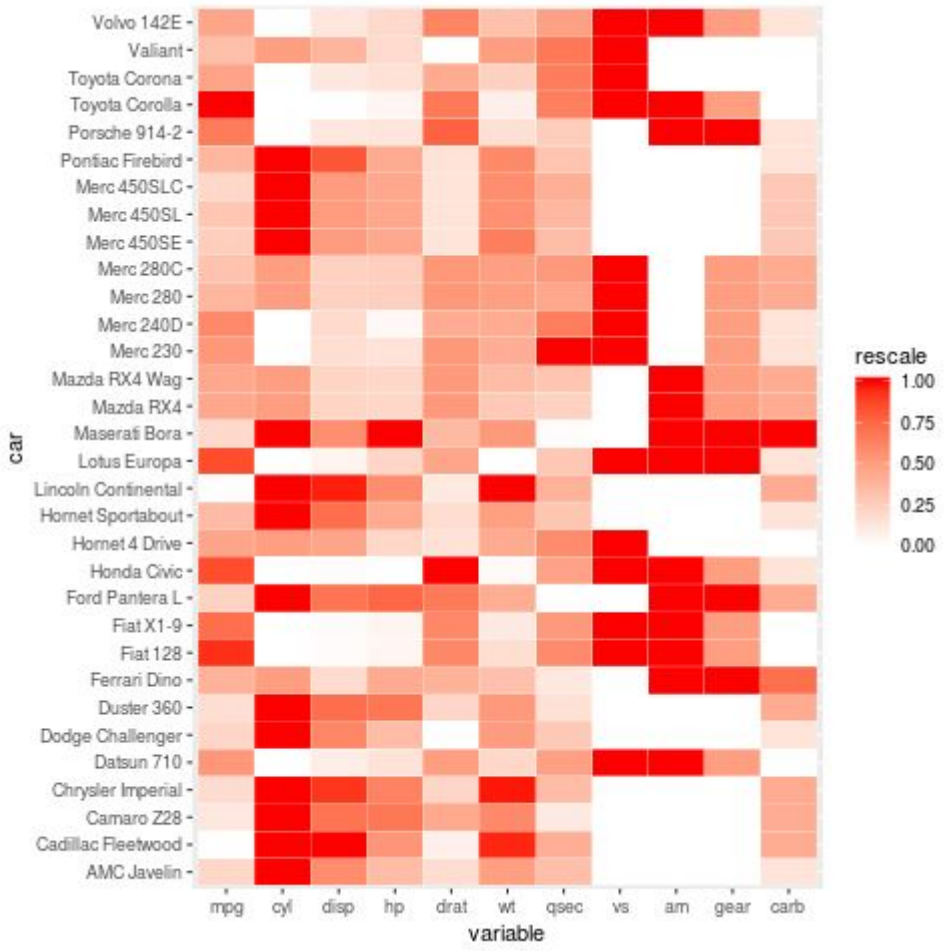


The resulting visualization is significantly improved. By using the `rescale` column for the fill aesthetic, the color gradient now accurately reflects the relative intensity of each metric for each vehicle, allowing for clear visual comparison across all variables in the **mtcars** dataset. This step is critical for any heatmap visualization involving variables measured on vastly different scales.

## Customizing Aesthetics: Color Scales and Ordering

Once the data is correctly scaled, we can focus on aesthetic refinements to improve the visual appeal and interpretability of the **heatmap**. A simple yet powerful customization is adjusting the color palette. While red is commonly used to indicate high intensity, choosing an appropriate color gradient can enhance clarity, especially when preparing visualizations for publication or presentation. For instance, a blue gradient often provides a cleaner, more neutral look.

Changing the color scheme requires only a small modification to the `scale_fill_gradient()` argument, replacing the `high` color parameter from "red" to "steelblue". This flexibility is a hallmark of the **ggplot2** package, allowing for rapid iteration on visual design without complex code changes. The code below demonstrates this change:

```
#create heatmap using blue color scale
ggplot(melt_mtcars, aes(variable, car)) +
geom_tile(aes(fill = rescale), colour = "white") +
scale_fill_gradient(low = "white", high = "steelblue")
```

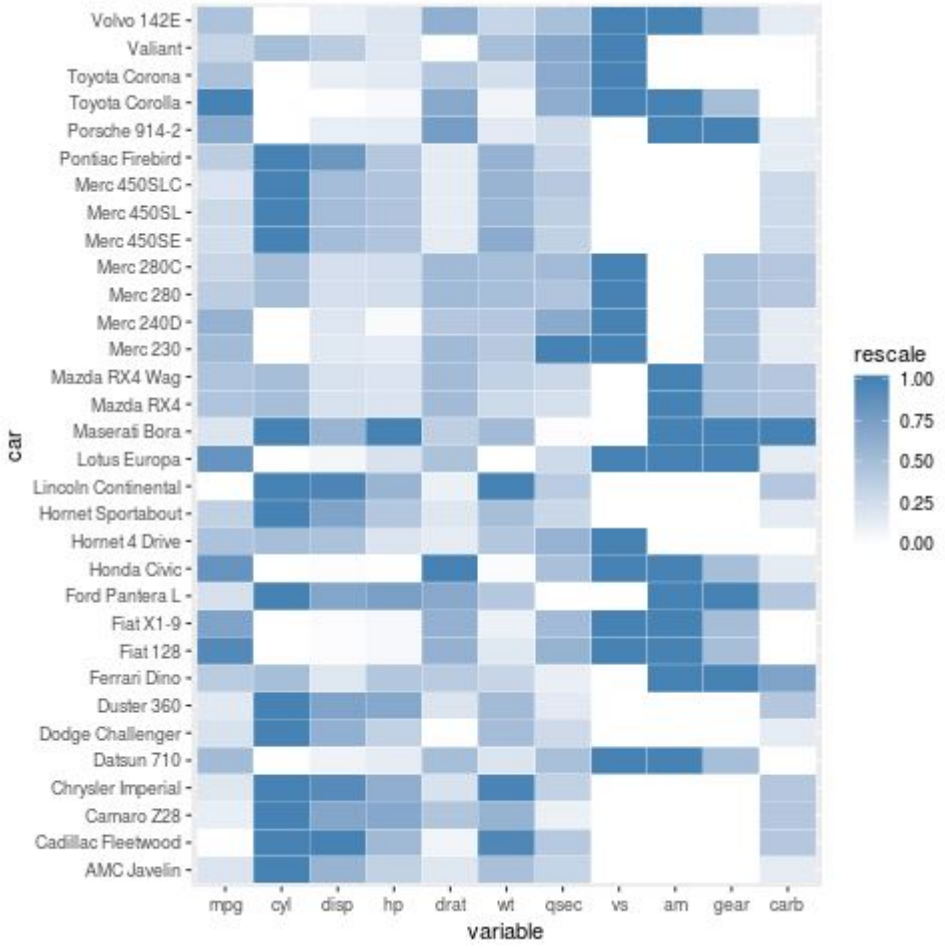


Another crucial aspect of heatmap design is the ordering of the observations (the y-axis rows). By default, the car names are ordered alphabetically, which often masks underlying patterns. To maximize the visual clustering and enhance interpretability, it is highly beneficial to order the rows based on the values of a meaningful variable. For example, we might want to group cars based on their fuel efficiency, measured by *mpg* (miles per gallon).

## Advanced Customization: Reordering Rows Based on Variable Values

To reorder the car names (rows) based on a specific variable like *mpg*, we must manipulate the factor levels of the `car` column in the original data frame before melting it. The `reorder()` function, often used in conjunction with `with()` for clarity, allows us to define the order of the factor levels

based on the corresponding numerical values in the *mpg* column. By ordering the data frame rows based on *mpg*, the heatmap will visually cluster similar performing vehicles together.

This reordering step must occur prior to the melting and rescaling steps, as the factor order established here is preserved when the data is transformed into the long format. The following comprehensive code chunk incorporates the reordering, melting, and rescaling steps to produce a heatmap where the cars are sorted by their *mpg* in ascending order (lowest mpg at the bottom, highest at the top):

```
#define car name as a new column, then order by mpg ascending
mtcars$car <- row.names(mtcars)
mtcars$car <- with(mtcars, reorder(car, mpg))

#melt mtcars into long format
melt_mtcars <- melt(mtcars)

#rescale values for all variables in melted data frame
melt_mtcars <- ddply(melt_mtcars, .(variable), transform, rescale = rescale(value))

#create heatmap using rescaled values
ggplot(melt_mtcars, aes(variable, car)) +
geom_tile(aes(fill = rescale), colour = "white") +
scale_fill_gradient(low = "white", high = "steelblue")
```

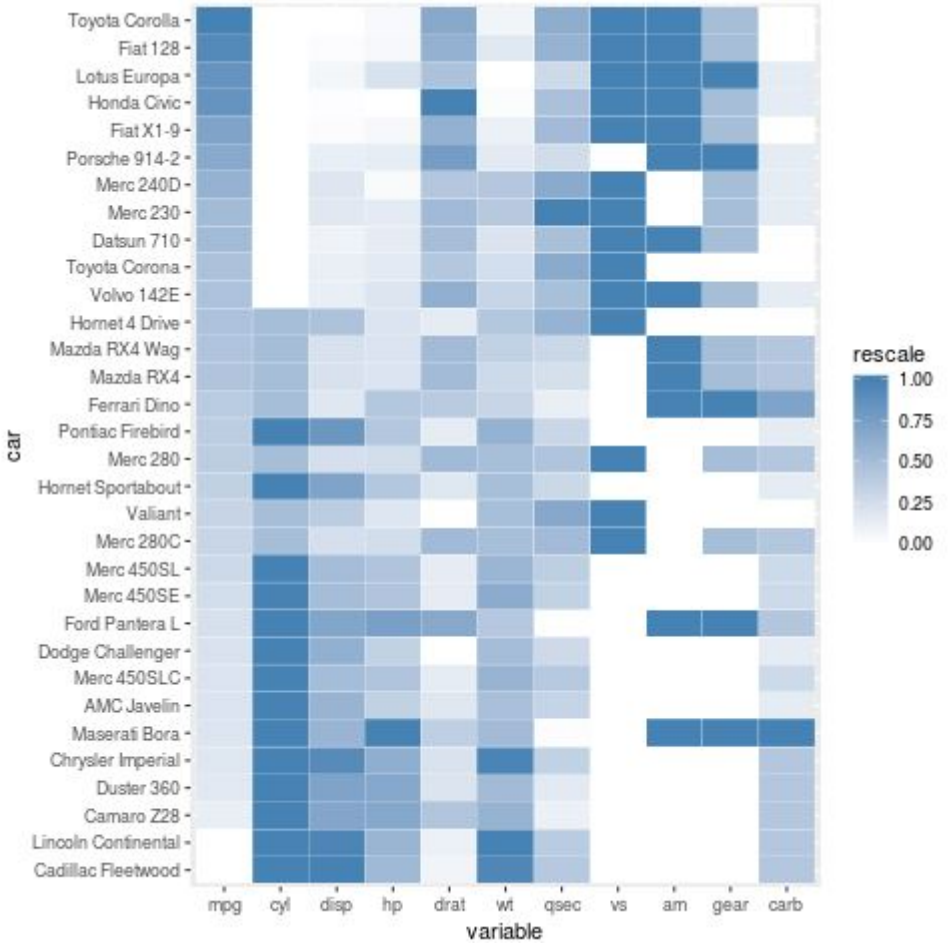


If the goal is to sort the car models in the opposite direction--by *mpg* descending--the modification is minimal but important. Within the `reorder()` function, we simply multiply the sorting variable by negative one (`-mpg`). This reverses the inherent ordering mechanism, placing the highest *mpg* values at the bottom of the visualization and the lowest values at the top. This technique provides complete control over the visual hierarchy presented in the final **heatmap**, directly impacting how the viewer interprets patterns of performance across the fleet of vehicles.

**#define car name as a new column, then order by mpg descending**
**mtcars$car <- row.names(mtcars)**
**mtcars$car <- with(mtcars, reorder(car, -mpg))**

#melt mtcars into long format
melt_mtcars <- melt(mtcars)

#rescale values for all variables in melted data frame
melt_mtcars <- ddply(melt_mtcars, .(variable), transform, rescale = rescale(value))

#create heatmap using rescaled values

```
ggplot(melt_mtcars, aes(variable, car)) +
geom_tile(aes(fill = rescale), colour = "white") +
scale_fill_gradient(low = "white", high = "steelblue")
```

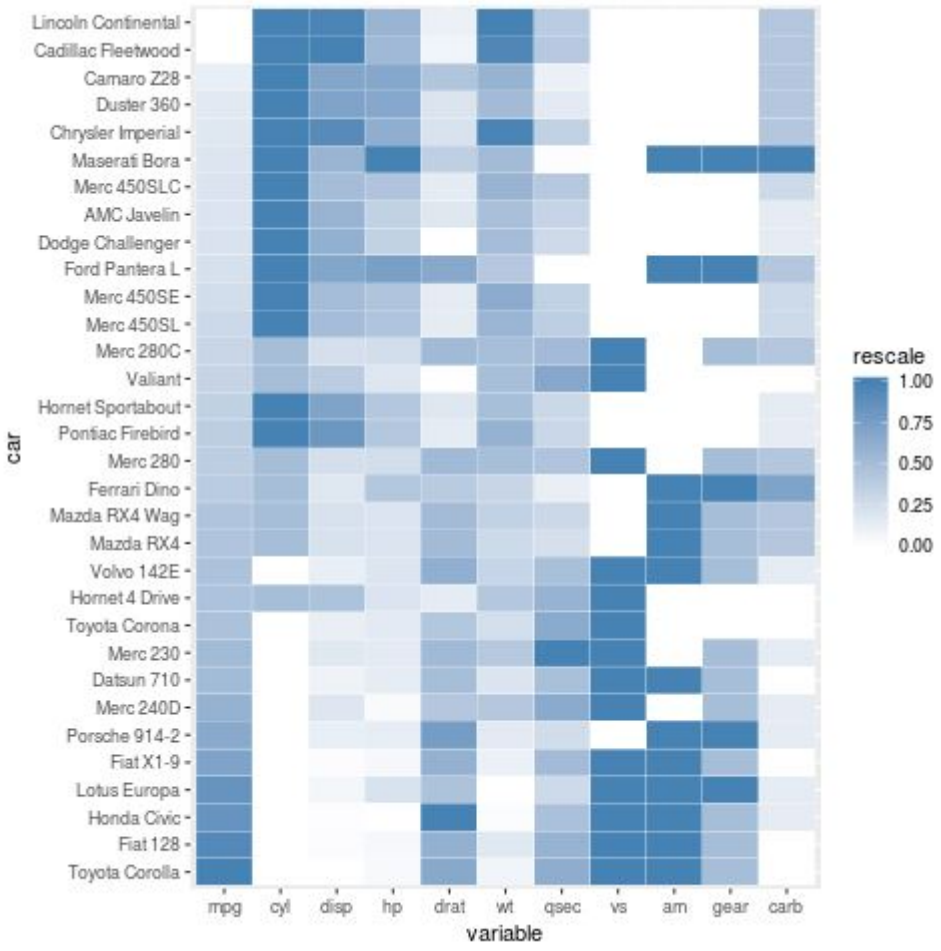


## Final Polish: Removing Axis Labels and Legends

In many presentation contexts, especially when the axes are self-explanatory or detailed annotations are provided elsewhere, removing extraneous elements like axis labels and the legend can clean up the visualization and allow the data itself to take center stage. The **ggplot2** package provides straightforward methods for these final adjustments using the `labs()` and `theme()` functions.

The `labs()` function is used to modify plot labels, including the x and y axes. By setting both `x` and `y` arguments to an empty string (`""`), we effectively remove the axis titles. Furthermore, the `theme()` function controls non-data elements of the plot. Setting `legend.position = "none"` ensures that the color key, which shows the relationship between color and the `rescale` value, is

hidden. This is often done when the relative comparisons within the heatmap are more important than the exact numerical values represented by the scale.

The final, polished code snippet demonstrates how to achieve this minimalist aesthetic:

```
#create heatmap with no axis labels or legend
ggplot(melt_mtcars, aes(variable, car)) +
geom_tile(aes(fill = rescale), colour = "white") +
scale_fill_gradient(low = "white", high = "steelblue") +
labs(x = "", y = "") +
theme(legend.position = "none")
```

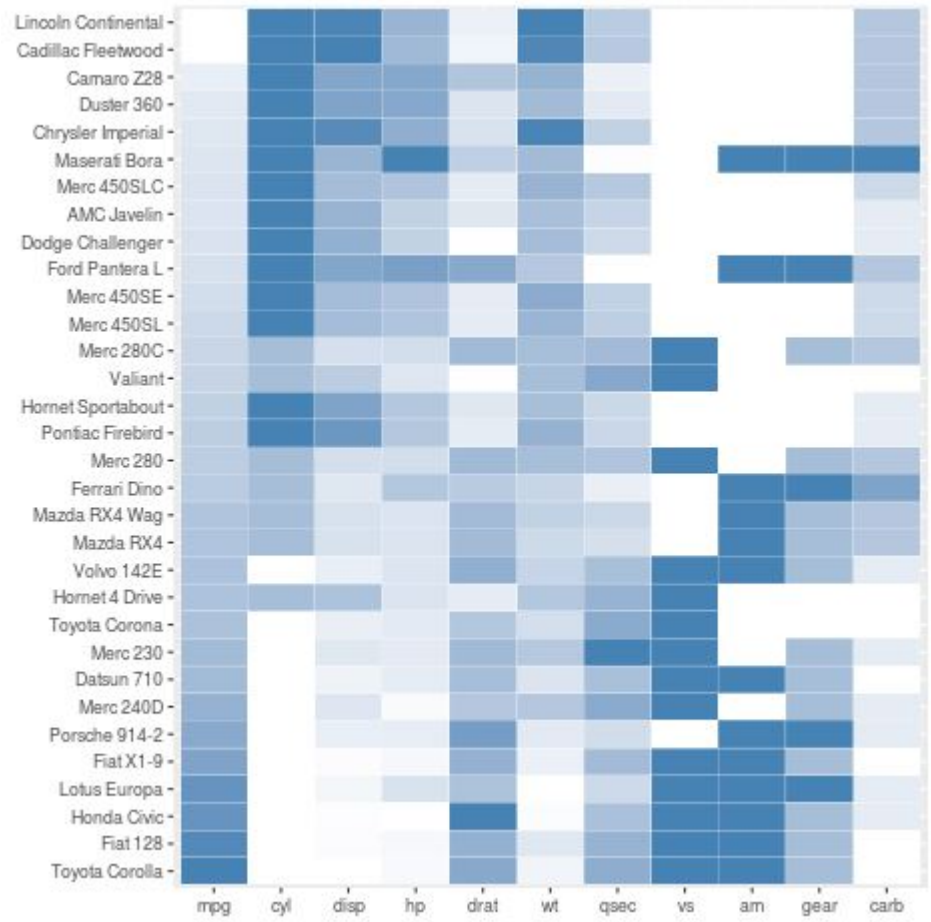


By following these steps--from data preparation and transformation to rescaling and advanced aesthetic customization--you can generate highly informative and visually appealing heatmaps using **R** and **ggplot2**. Heatmaps remain an indispensable tool for multivariate data exploration, offering immediate insight into relationships and patterns that might be obscured in raw data tables.