

How to Plot Multiple Linear Regression Results in R?

Authored by
stats writer

December 14, 2025

RECOMMENDED CITATION

stats writer (2025). *How to Plot Multiple Linear Regression Results in R?*.
PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=107464>

Visualizing the results of a Multiple Linear Regression (MLR) model in R requires specialized techniques beyond simple 2D plotting. While simple linear regression allows for straightforward plotting using functions like the lm() command to define the model, the plot() command to generate a scatterplot, and the abline() command to superimpose the regression line, MLR introduces complexities due to multiple predictor variables. Effective visualization of MLR often involves generating plots that isolate the effect of each predictor, such as those that display confidence intervals, residuals, and predicted values across dimensions.

When conducting a simple linear regression analysis in R, the process of visualizing the fitted model is intuitively simple. Since we are typically working with only one predictor and one response variable, the relationship exists within a two-dimensional space. This allows us to easily map the data points and draw the best-fit line that minimizes the sum of squared errors.

For instance, the following R code demonstrates how to fit a basic simple linear regression model to a sample dataset and subsequently plot the empirical data points along with the estimated regression line. This visualization is fundamental for assessing the relationship's direction and magnitude in a simple context.

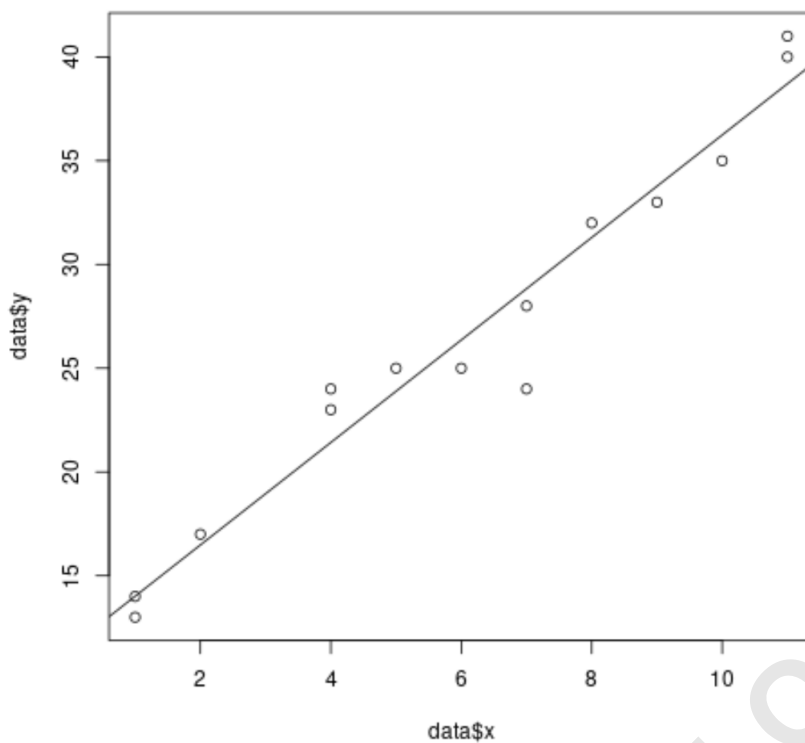
Create a sample dataset for demonstration

```
data <- data.frame(x = c(1, 1, 2, 4, 4, 5, 6, 7, 7, 8, 9, 10, 11, 11),  
y = c(13, 14, 17, 23, 24, 25, 25, 24, 28, 32, 33, 35, 40, 41))
```

```
# Fit a simple linear regression model using the lm() command  
model <- lm(y ~ x, data = data)
```

```
# Create a scatterplot of the predictor variable (x) against the response variable (y)  
plot(data$x, data$y)
```

```
# Add the fitted regression line to the scatterplot using the abline() command  
abline(model)
```



The Visualization Challenge in Multiple Linear Regression

When the analysis shifts to Multiple Linear Regression, the inherent ease of visualization disappears. A multiple regression model involves two or more predictor variables used to explain variation in a single response variable. For instance, a model with three predictors exists in a four-dimensional space, which cannot be accurately or fully represented on a conventional two-dimensional plot or even a three-dimensional surface plot.

The core challenge is that the estimated coefficient for any single predictor variable in an MLR model reflects its relationship with the response variable while statistically holding all other predictor variables constant. This concept of "controlling for" or "partial effect" is essential but difficult to illustrate in a standard scatterplot of the raw data. Plotting the raw response variable against a single predictor variable, without accounting for the presence of the others, often yields a misleading representation of that variable's true contribution to the model.

Introducing Added Variable Plots (Partial Regression Plots)

To overcome the dimensional limitations of Multiple Linear Regression visualization, statisticians rely on diagnostic plots designed to display these critical partial effects. The most common and effective tool for this is the use of added variable plots, often referred to as partial regression plots. These individual plots display the adjusted relationship between the response variable and one

specific predictor variable, effectively controlling for the influence of every other predictor variable included in the statistical model.

An added variable plot for a predictor variable (X_k) is constructed by plotting two sets of residuals against each other. The vertical axis plots the residuals from the regression of the response variable (Y) on all other predictor variables (excluding X_k). The horizontal axis plots the residuals from the regression of the specific predictor variable (X_k) on all other predictor variables (also excluding X_k). The resulting linear fit in this plot provides a clear, two-dimensional view of the unique contribution of X_k to explaining the variation in Y .

The following detailed example demonstrates how to execute a Multiple Linear Regression analysis in R and subsequently visualize the results using added variable plots to isolate the effect of each regressor.

Setting Up the Multiple Linear Regression Model Example

For this demonstration, we will utilize the built-in R dataset, **mtcars**, which contains data on various characteristics of 32 automobiles. We aim to model the fuel efficiency, measured in miles per gallon (mpg), as a function of three predictors: engine displacement (disp), horsepower (hp), and rear axle ratio (drat).

We use the lm() command to fit the model and the R summary() function to examine the initial statistical output, including coefficient estimates, standard errors, t-values, and p-values. This initial output helps confirm the overall fit and significance of the model before proceeding to graphical diagnostics.

```
# Fit the multiple linear regression model: mpg explained by disp, hp, and drat  
model <- lm(mpg ~ disp + hp + drat, data = mtcars)
```

```
# View the detailed results of the model fitting  
summary(model)
```

Call:

```
lm(formula = mpg ~ disp + hp + drat, data = mtcars)
```

Residuals:

```
Min 1Q Median 3Q Max
```

```
-5.1225 -1.8454 -0.4456 1.1342 6.4958
```

Coefficients:

```
Estimate Std. Error t value Pr(>|t|)
```

```
(Intercept) 19.344293 6.370882 3.036 0.00513 **
```

```
disp -0.019232 0.009371 -2.052 0.04960 *
hp -0.031229 0.013345 -2.340 0.02663 *
drat 2.714975 1.487366 1.825 0.07863 .
---
Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

```
Residual standard error: 3.008 on 28 degrees of freedom
Multiple R-squared: 0.775, Adjusted R-squared: 0.7509
F-statistic: 32.15 on 3 and 28 DF, p-value: 3.28e-09
```

Upon reviewing the model summary, we observe that the overall model is highly significant (p -value $< 3.28e-09$). Furthermore, the individual p -values for the coefficients of `disp` and `hp` are below 0.05, suggesting their statistical significance. The coefficient for `drat` is also marginally significant ($p=0.07863$). For the purposes of visualization, we proceed assuming all these variables are relevant and should remain in the model, justifying the need to visualize their partial effects.

Implementing Added Variable Plots using the 'car' Package

To generate the required added variable plots in R, we rely on the powerful `car` package (Companion to Applied Regression), developed by John Fox and Sanford Weisberg. This package provides robust tools for regression diagnostics, including the `avPlots()` function, which automates the creation of these specialized residual plots.

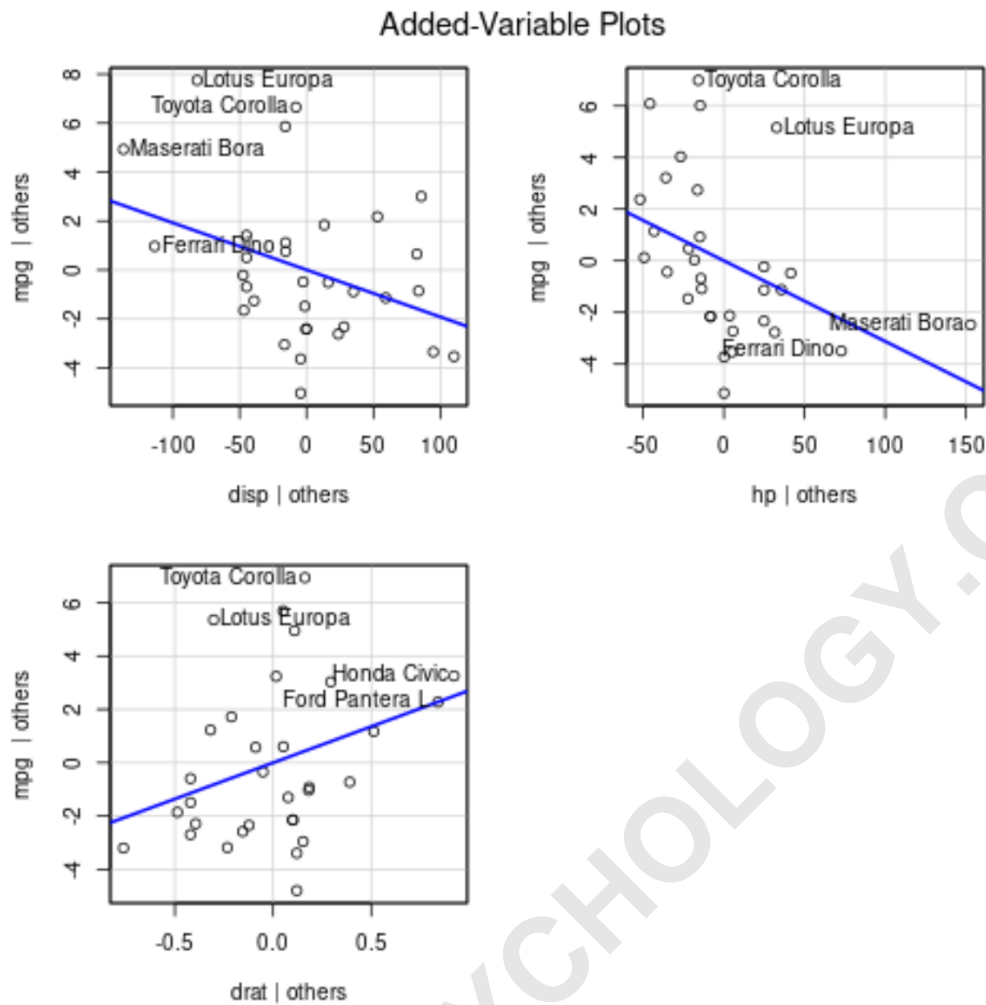
The `avPlots()` function takes the fitted `lm()` model object as its primary argument and automatically generates a separate partial scatterplot for each predictor variable within the model. This is the simplest and most efficient way to visualize the partial linear relationships in an MLR context.

Load the 'car' package, which contains the avPlots() function

```
library(car)
```

```
# Produce the added variable plots for the fitted multiple regression model
```

```
avPlots(model)
```



Detailed Interpretation of Added Variable Plots

Each panel in the generated output corresponds to a single predictor variable (disp, hp, or drat). Understanding the axes and the elements within these plots is crucial for effective diagnosis of the MLR model:

The **x-axis** displays the residuals of the predictor variable in question, regressed against all other predictors in the model. This represents the part of the predictor variable that is independent of the other regressors.

The **y-axis** displays the residuals of the response variable (mpg), regressed against all other predictor variables. This represents the unexplained variation in the response variable after accounting for the influence of every other factor.

The **Blue Line** (the partial regression line) illustrates the estimated partial relationship between the respective predictor variable and the response variable. The slope of this line is precisely the

estimated coefficient for that predictor in the full Multiple Linear Regression model. This line shows the association while statistically holding the value of all other predictor variables constant.

The **Labeled Points** represent influential observations within the dataset. Typically, the plots identify the two data points with the largest partial leverage and the two data points with the largest Studentized residuals. Identifying these influential points is vital for determining if specific observations unduly bias the coefficient estimates or standard errors.

Connecting Plot Angles to Regression Coefficients

One of the most valuable aspects of using added variable plots is the direct visual confirmation of the sign and magnitude of the estimated regression coefficients. The angle, or slope, of the partial regression line in each plot directly matches the numerical coefficient found in the output of the summary() function.

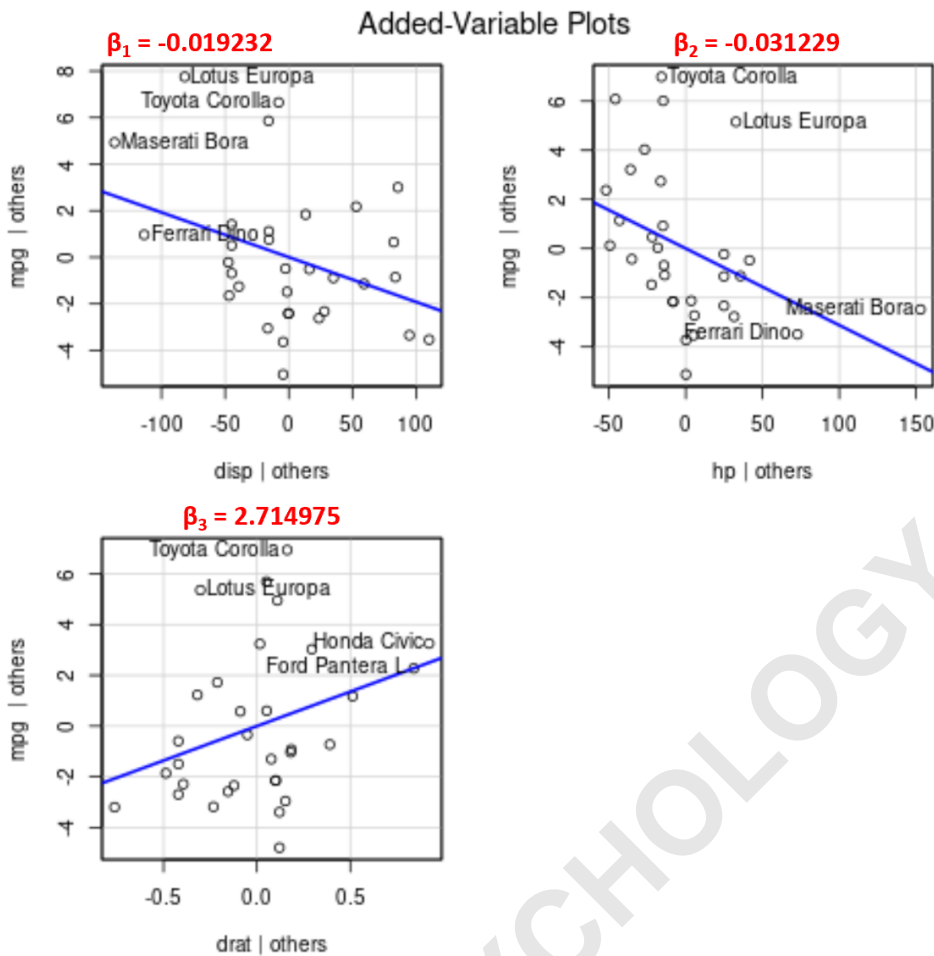
Let's revisit the coefficients derived from our Multiple Linear Regression model for the mtcars dataset:

disp: -0.019232 (Negative coefficient)

hp: -0.031229 (Negative coefficient)

drat: 2.714975 (Positive coefficient)

The visualization perfectly aligns with these numerical results. The partial regression line for **drat** exhibits a clear positive slope, indicating that, holding displacement and horsepower constant, a higher rear axle ratio is associated with increased fuel efficiency (mpg). Conversely, the lines for **disp** and **hp** both angle downwards, confirming their negative estimated coefficients: increases in engine displacement or horsepower are associated with decreased fuel efficiency, even after accounting for the other variables in the model.



Conclusion on Visualizing Partial Effects

Although the inherent dimensionality of Multiple Linear Regression prevents the plotting of a single, encompassing fitted regression line on a two-dimensional surface, added variable plots offer an essential diagnostic tool. These plots successfully decompose the complex relationships, allowing analysts to observe the unique, partial relationship between each individual predictor variable and the response variable.

This visualization method is critical for several statistical purposes, including: verifying the linearity assumption for each predictor, detecting influential points that might distort coefficient estimates, and providing strong visual evidence supporting the direction (sign) and strength of the partial correlations identified in the numerical regression output. By utilizing the **car** package and the **avPlots()** function in R, advanced regression models can be interpreted with clarity and confidence.