

Learning to Visualize Interactions: A Guide to Creating Interaction Plots in R for Two-Way ANOVA

Authored by
Mohammed loot

November 7, 2025

RECOMMENDED CITATION

Mohammed loot (2025). *Learning to Visualize Interactions: A Guide to Creating Interaction Plots in R for Two-Way ANOVA*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=12184>

Understanding Interaction Effects in Statistical Modeling

The [two-way ANOVA](#) is a powerful statistical technique utilized to assess whether the means of a continuous outcome variable differ across groups defined by two distinct categorical factors. This method allows researchers to simultaneously evaluate the independent effects of each factor, known as **main effects**, and the joint effect of both factors working together.

We employ a [two-way ANOVA](#) when our objective is to determine how two specific independent variables influence a single [response variable](#). This is critical in experimental design where control over multiple factors is necessary for accurate interpretation of results. However, a complexity often arises when the effect of one factor is dependent on the level of the other factor; this phenomenon is referred to as an **interaction effect**.

The presence of a significant [interaction effect](#) can profoundly impact how we interpret the relationship between the factors and the [response variable](#). If an interaction exists, the main effects alone may provide a misleading or incomplete picture. For instance, consider an experiment investigating whether the factors **exercise intensity** and **gender** affect the [response variable](#) of **weight loss**. While both factors might independently contribute to weight loss, it is also highly plausible that the rate of weight loss due to intense exercise differs significantly between males and females. In this specific scenario, we would conclude that there is an interaction between exercise and gender, meaning the effect of exercise is moderated by gender.


The Role and Interpretation of an Interaction Plot

To properly detect, visualize, and understand these complex interactions, statisticians rely on graphical tools, most notably the **interaction plot**. This visualization is essential because while statistical tests like ANOVA can confirm the existence of an [interaction effect](#), the plot illustrates the nature and magnitude of that interaction across all factor levels, providing critical context for interpretation.

An interaction plot systematically displays the fitted values (or means) of the [response variable](#) on the y-axis. The values of the first categorical factor are plotted along the x-axis. Crucially, the levels of the second factor are represented by distinct lines (or traces) drawn across the plot. By observing the slopes and parallelism of these lines, we gain immediate insight into how the factors combine to influence the outcome.

The primary rule for interpreting an interaction plot is straightforward: if the lines representing the second factor are generally **parallel**, there is typically no significant [interaction effect](#). This suggests that the effect of the first factor is consistent across all levels of the second factor. Conversely, if the lines exhibit a lack of parallelism--meaning they cross, converge, or diverge significantly--this strongly suggests the presence of a meaningful interaction. Intersecting lines are

the clearest visual indicator that the effect of one factor depends directly on the level of the other factor.



Case Study: Impact of Exercise and Gender on Weight Loss

Consider a hypothetical research study designed to rigorously examine the relationship between physical activity and biological sex on weight management outcomes. The researchers hypothesize that both **exercise intensity** (Factor 1) and **gender** (Factor 2) influence the outcome of **weight loss** (the continuous response variable). To test this, they recruit 60 participants: 30 men and 30 women.

The participants are randomly assigned to one of three exercise programs for a one-month duration: **No Exercise**, **Light Exercise**, or **Intense Exercise**. This design ensures that ten men and ten women are allocated to each of the three intensity levels. The resulting dataset contains 60 observations, structured to facilitate a [two-way ANOVA](#), which will analyze the main effects of exercise and gender, as well as their potential interaction.

The critical question driving this analysis is whether the beneficial effects of increasing exercise intensity on weight loss are consistent regardless of the participant's gender. If, for example, women experience substantially greater weight loss improvements when moving from light to intense exercise compared to men, then the interaction term will be statistically significant. This necessitates the use of the interaction plot to visualize exactly where and how these differences

manifest across the treatment levels.

Step-by-Step Implementation in R

The process of analyzing this experimental data begins with preparing the data structure within the R environment, fitting the appropriate statistical model, and finally, generating the visualization. We begin by constructing the [data frame](#) that accurately represents the experimental conditions and the observed outcomes.

Step 1: Creating the Data Frame

The following R code snippet demonstrates how to generate the synthetic data set, ensuring the experimental design (30 males, 30 females, 10 participants per exercise level within each gender group) is accurately reflected. The `set.seed()` function is used to ensure the results are reproducible, a standard practice in statistical programming. The `runif()` function simulates the weight loss values based on predefined ranges that reflect the expected outcomes for each group.

#make this example reproducible

set.seed(10)

#create data frame

```
data <- data.frame(gender = rep(c("Male", "Female"), each = 30),  
  exercise = rep(c("None", "Light", "Intense"), each = 10, times = 2),  
  weight_loss = c(runif(10, -3, 3), runif(10, 0, 5), runif(10, 5, 9),  
    runif(10, -4, 2), runif(10, 0, 3), runif(10, 3, 8)))
```

#view first six rows of data frame

head(data)

gender exercise weight_loss

1 Male None 0.04486922

2 Male None -1.15938896

3 Male None -0.43855400

4 Male None 1.15861249

5 Male None -2.48918419

6 Male None -1.64738030

Step 2: Fitting the Two-Way ANOVA Model

Once the data is prepared, the next step involves fitting the two-way ANOVA model using the `aov()` function in R. The formula `weight_loss ~ gender * exercise` specifies that we are modeling

the `weight_loss` outcome based on the main effects of `gender` and `exercise`, plus their multiplicative interaction term (`gender * exercise`). This interaction term is the core focus of the analysis, as it tests whether the combination of factors yields a result different from the sum of their individual effects.

Analyzing the model summary is crucial for drawing statistical conclusions. We must pay close attention to the row labeled `gender:exercise` in the output table. This row provides the F-statistic and the corresponding [p-value](#) for the interaction hypothesis test. A small [p-value](#) (typically below 0.05) indicates that the interaction effect is statistically [significant](#), confirming that the effect of exercise indeed varies depending on gender.

#fit the two-way ANOVA model

```
model <- aov(weight_loss ~ gender * exercise, data = data)
```

```
#view the model output
```

```
summary(model)
```

```
# Df Sum Sq Mean Sq F value Pr(>F)
```

```
#gender 1 15.8 15.80 11.197 0.0015 **
```

```
#exercise 2 505.6 252.78 179.087 <2e-16 ***
```

```
#gender:exercise 2 13.0 6.51 4.615 0.0141 *
```

```
#Residuals 54 76.2 1.41
```

```
#---
```

```
#Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

In this specific output, the [p-value](#) for the interaction term (`gender:exercise`) is **0.0141**. Since this value is less than the conventional alpha level of 0.05, we conclude that the interaction between exercise intensity and gender is statistically [significant](#). This result confirms the necessity of moving beyond the summary table and visualizing the effect, as the interaction term complicates the interpretation of the main effects.

Visualizing the Interaction: Generating the Plot in R


The third step involves generating the interaction plot using R's built-in `interaction.plot()` function. This function provides a quick and effective visual representation of the group means, allowing for a descriptive assessment of the interaction detected by the [two-way ANOVA](#). Proper specification of the function arguments is vital to ensure the plot accurately reflects the model structure.

When calling `interaction.plot()`, we define the `x.factor` (the variable plotted on the horizontal axis, in this case, exercise), the `trace.factor` (the variable defining the separate lines, here gender), and the `response` variable (weight loss). We specify `fun = median` to plot the median weight loss for

each group combination, a choice that can be more robust to outliers than plotting the mean, although the mean is also commonly used. Customization arguments such as ``ylab``, ``xlab``, ``col``, and ``trace.label`` enhance the clarity and professionalism of the resulting visualization.

```
interaction.plot(x.factor = data$exercise, #x-axis variable
trace.factor = data$gender, #variable for lines
response = data$weight_loss, #y-axis variable
fun = median, #metric to plot
ylab = "Weight Loss",
xlab = "Exercise Intensity",
col = c("pink", "blue"),
lty = 1, #line type
lwd = 2, #line width
trace.label = "Gender")
```

The execution of this code yields the interaction plot, which is the necessary complement to the numerical output of the ANOVA. It graphically summarizes the conditional effects, enabling researchers to pinpoint precisely which levels of exercise intensity produce different results for males versus females.



Interpreting the Interaction Plot Results

As established, the visual interpretation hinges on the parallelism of the lines. If the effect of exercise on weight loss were identical for both genders, the pink line (Females) and the blue line (Males) would run roughly parallel to each other across all three levels of exercise intensity (None, Light, Intense). However, the plot generated by the synthetic data clearly shows a non-parallel pattern.

A close inspection of the visualization reveals that the two lines intersect between the "Light" and "Intense" exercise categories. This intersection is the hallmark of a strong [interaction effect](#). Specifically, the plot illustrates that for participants engaged in "No Exercise," males and females show similar (low) weight loss medians. At the "Light Exercise" level, females show a slightly higher median weight loss than males. Crucially, at the "Intense Exercise" level, the trend reverses sharply, with males achieving significantly higher weight loss than females.

This visual evidence confirms the statistical finding from the previous step: the effect of exercise intensity on weight loss is conditional upon the participant's gender. The statistically [significant p-value](#) (0.0141) from the ANOVA output for the interaction term is fully supported by the non-parallel and intersecting lines observed in the interaction plot. Consequently, researchers must interpret the results not by looking at gender or exercise in isolation, but by analyzing the specific combinations of these two factors.

Additional Resources

[How to Conduct a One-Way ANOVA in R](#)

[How to Conduct a Two-Way ANOVA in R](#)