

How to create an Interaction Plot in R

Authored by
stats writer

December 21, 2025

RECOMMENDED CITATION

stats writer (2025). *How to create an Interaction Plot in R*. PSYCHOLOGICAL SCALES.
Retrieved from <https://scales.arabpsychology.com/?p=108258>

The Interaction Plot is an essential statistical tool used to visualize how the relationship between a response variable and one explanatory variable changes across the levels of a second categorical variable. In R, generating this visualization is straightforward using the built-in `interaction.plot()` function. This function requires three key inputs: the variable being measured (the response variable), the primary factor displayed on the x-axis (the explanatory variable), and the secondary factor represented by different lines (the interaction variable).

Understanding how to utilize `interaction.plot()` is critical for researchers analyzing complex experimental designs, particularly those involving factorial Analysis of Variance (ANOVA) models. The plot translates complex numerical relationships into a clear graphical format, positioning the mean (or median) of the response variable on the y-axis, mapped against the levels of the two interacting factors on the x-axis and via distinct traces. These visual representations are indispensable for interpreting sophisticated statistical models accurately.

Introduction to Interaction Plots and ANOVA


When analyzing data where a single outcome is influenced by two categorical independent variables, the two-way ANOVA (Analysis of Variance) is the standard statistical test. This technique determines whether the means of multiple groups differ significantly, based on the stratification provided by two distinct factors. For instance, a researcher might use this approach to evaluate how Factors A and B individually influence a specific outcome measurement, known as the response variable.

While the primary focus of ANOVA often lies in assessing the main effects of each factor, a crucial consideration is the potential presence of an interaction effect. An interaction occurs when the effect of one factor on the response variable depends on the level of the other factor. Ignoring a significant interaction can lead to misleading conclusions about the main effects, rendering their interpretation incomplete or even incorrect.

Consider a practical example: evaluating the impact of exercise intensity (Factor 1) and gender (Factor 2) on weight loss (the response variable). While both exercise and gender might independently influence weight loss, an interaction effect suggests that the benefit derived from a specific exercise intensity level (e.g., intense exercise) is substantially different for males compared to females. It is possible that intense exercise yields dramatically higher weight loss for one gender, but only marginally higher loss for the other.

Visualizing this complex relationship is best achieved using an interaction plot. This visualization specifically displays the mean values of the response variable (Y-axis) relative to the levels of the first factor (X-axis), with separate lines (traces) dedicated to representing each level of the second factor. The pattern and slope of these lines immediately reveal the nature and magnitude of the

interaction between the two independent variables.



This detailed guide explains how to properly set up, execute, and interpret an interaction plot in R, confirming the results against a formal statistical test.

Interpreting Interaction Plots: Parallel vs. Intersecting Lines

The visual analysis of an interaction plot is based on observing the relationship between the lines representing the different groups (levels of the second factor). This visual interpretation provides immediate insight into the statistical significance and practical relevance of the observed interaction effect. The two primary patterns to watch for are parallel lines and non-parallel (intersecting or converging/diverging) lines.

If the lines corresponding to the different levels of the trace variable are approximately **parallel** across the x-axis, it suggests that the effect of the primary factor (x-axis) is consistent across all levels of the secondary factor. Statistically, parallel lines indicate a lack of a significant interaction effect, meaning that the factors operate independently on the response variable. In such a scenario, the main effects of the ANOVA model can generally be interpreted without qualification.

Conversely, if the lines exhibit patterns of **intersection**, convergence, or divergence, it is a strong visual indicator of a significant interaction effect. Intersecting lines imply that the effect of the primary factor changes direction or magnitude depending on which level of the secondary factor is

being considered. For instance, if one line slopes steeply upward while another remains flat or slopes downward, it confirms that the factors are interdependent in their impact on the response variable. This visual evidence should always be validated by the corresponding statistical significance reported in the ANOVA output, specifically the p-value for the interaction term.

Detailed Components of the `interaction.plot()` Function

The core function used in R for generating these visualizations is `interaction.plot()`. To ensure the plot accurately reflects the intended experimental design, it is essential to correctly specify the function's arguments. Unlike plotting functions that automatically infer roles, this function requires explicit assignment of the three main variable types.

The primary arguments required for successful execution include:

`x.factor`: This specifies the categorical variable that will be displayed along the horizontal axis. In a factorial design, this is typically one of the two independent factors, such as "Exercise Intensity."

`trace.factor`: This defines the second categorical independent variable, whose levels are represented by distinct lines (traces) on the plot. For example, "Gender" might be assigned to this argument.

`response`: This is the quantitative dependent variable, whose aggregated values (means or medians) are plotted on the vertical axis (Y-axis). In our example, this would be "Weight Loss."

Furthermore, the function offers powerful optional arguments that allow for customization and clarity:

`fun`: Specifies the function used to aggregate the response variable values at each combination of factor levels. By default, this is the `mean`, but using `median` is often recommended if the data contains outliers or is not normally distributed.

`ylab` and `xlab`: Used to provide clear, descriptive labels for the Y and X axes, respectively, greatly enhancing the plot's readability.

`col`, `lty`, and `lwd`: Control the visual aesthetics of the plot, including the color of the lines (`col`), the line type (`lty`, e.g., solid or dashed), and the line width (`lwd`). Customizing these parameters ensures differentiation between the trace factors.

Case Study: Evaluating Factors Influencing Weight Loss

To illustrate the complete process of detecting and visualizing interaction effects, we will explore a common scenario in health research. Suppose a team of researchers aims to rigorously determine

whether the intervention factors of **exercise intensity** and **gender** jointly influence the outcome of **weight loss** over a standardized period. This experiment follows a classic two-way ANOVA design.

The experimental setup involves recruiting 60 participants: 30 men and 30 women. These participants are then randomly assigned 10 of each gender to one of three exercise programs: **None**, **Light exercise**, or **Intense exercise**. The study duration is one month, after which the net weight loss (in kilograms) is recorded for every individual. This design allows for the calculation of the main effects of exercise and gender, as well as their crucial interaction effect.

The following steps outline the necessary procedures within R: first, generating a synthetic data structure that mirrors the experimental outcomes; second, fitting the statistical ANOVA model to test for significance; and finally, generating the interaction plot to visually confirm and interpret the findings derived from the hypothesis test between exercise and gender.

Step 1: Preparing the Data Frame in R

The initial phase involves constructing a suitable data structure to hold the experimental results. In R, this is typically a data frame, which organizes the data into columns representing variables and rows representing individual observations. For reproducibility and consistent results, we start by setting a seed value.

The data frame must include three columns corresponding to our factors and response variable: `gender` (a nominal factor), `exercise` (an ordinal factor), and `weight_loss` (a continuous numerical variable). The code below uses random uniform distribution functions (`runif`) to simulate realistic weight loss data corresponding to the three experimental conditions for each gender group, aligning with the 10 observations per subgroup established in the study design.

The following code snippet demonstrates the creation and structure verification of the necessary data frame:

```
#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 and Interpreting Results

After successfully preparing the data, the next critical step is to fit the two-way ANOVA model using the `aov()` function in R. This function tests the null hypothesis that the population means are equal across the levels of the factors. The model formula specifies `weight_loss` as the response variable, and `gender * exercise` indicates that we are interested in the main effects of gender and exercise, as well as their multiplicative interaction effect.

Fitting the model and examining the summary output provides the statistical evidence regarding the significance of the main and interaction effects. The Summary of the ANOVA table is structured to present the degrees of freedom (Df), sum of squares (Sum Sq), mean squares (Mean Sq), the F statistic (F value), and the critical p-value (Pr(>F)) for each tested term.

```
#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
```

Interpreting this output is crucial. We observe that both `gender` and `exercise` show statistically significant main effects (very low p-values). Most importantly, the row labeled `gender:exercise` represents the interaction term. With a p-value of **0.0141**, which is less than the conventional alpha level of 0.05, we confirm that the interaction effect between exercise intensity and gender is statistically significant. This significance confirms that the effect of exercise on weight loss is not

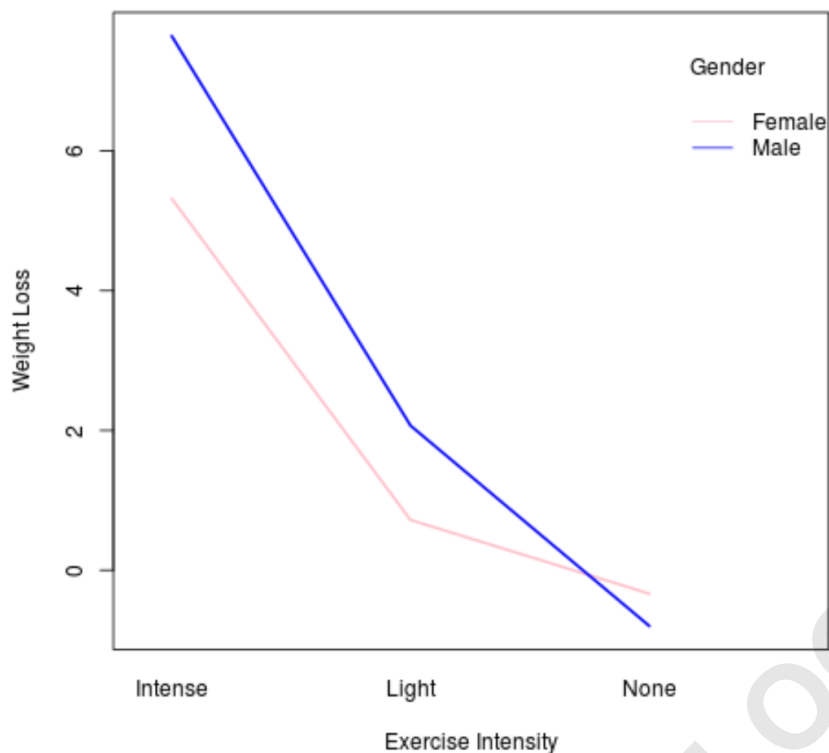
constant across both gender groups, necessitating a detailed visual examination using the interaction plot.

Step 3: Visualizing the Interaction Effect

The final step involves translating the statistical significance found in the ANOVA table into an intuitive visual format using the `interaction.plot()` function. Since the ANOVA demonstrated a significant interaction, the plot is necessary to understand the specific nature of this effect--that is, where and how the groups differ.

The following code generates the plot, explicitly defining `data$exercise` as the primary x-axis factor, `data$gender` as the trace factor (creating separate lines), and `data$weight_loss` as the response variable. We use `fun = median` for a robust measure of central tendency and apply customized labels and colors to ensure maximum clarity and professionalism in the final output.

```
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")
```



In general, if the two lines on the interaction plot are parallel, then there is no interaction effect. However, if the lines intersect, converge, or diverge significantly, an interaction effect is highly probable.

We can clearly see in this plot that the lines representing the median weight loss for males and females intersect between the "Light" and "Intense" exercise conditions. This distinct non-parallel pattern strongly indicates a significant interaction effect between the variables of exercise intensity and gender, confirming the statistical result ($p = 0.0141$) obtained from the ANOVA model.

This matches the fact that the p-value in the output of the ANOVA table was statistically significant for the interaction term in the ANOVA model.

[A Comprehensive Guide to Conducting a One-Way ANOVA in R](#)
[Mastering the Execution of a Two-Way ANOVA in R](#)