

How can I perform Tukey's test in R?

Authored by
stats writer

April 19, 2024

RECOMMENDED CITATION

stats writer (2024). *How can I perform Tukey's test in R?*. PSYCHOLOGICAL SCALES.
Retrieved from <https://scales.arabpsychology.com/?p=136991>

Tukey's test is a statistical method used to compare the means of multiple groups. In order to perform Tukey's test in R, you can follow these steps:

1. Load the data set into R: First, you need to import the data set that contains the groups you want to compare.
2. Install and load the "agricolae" package: Tukey's test is not built into R, so you will need to install the "agricolae" package. This package contains the necessary functions to perform Tukey's test.
3. Use the HSD.test" function: Once the package is loaded, you can use the HSD.test" function to perform Tukey's test. This function takes the data set as its input and calculates the confidence intervals and p-values for all possible pairwise comparisons between the groups.
4. Interpret the results: The output of the HSD.test" function will provide you with the confidence intervals and p-values for each comparison. You can use these results to determine if there are significant differences between the means of the groups.

In summary, to perform Tukey's test in R, you need to import the data set, install and load the "agricolae" package, use the HSD.test" function, and interpret the results to compare the means of multiple groups.

Perform Tukey's Test in R

A one-way ANOVA is used to determine whether or not there is a statistically significant difference between the means of three or more independent groups.

If the overall p-value from the ANOVA table is less than some significance level, then we have sufficient evidence to say that at least one of the means of the groups is different from the others.

However, this doesn't tell us *which* groups are different

from each other. It simply tells us that not all of the group means are equal. In order to find out exactly which groups are different from each other, we must conduct a post hoc test.

One of the most commonly used post hoc tests is Tukey's Test, which allows us to make pairwise comparisons between the means of each group while controlling for the family-wise error rate.

This tutorial explains how to perform Tukey's Test in R.

Note: If one of the groups in your study is considered a control group, you should instead use Dunnett's Test as the post-hoc test.

Example: Tukey's Test in R

Step 1: Fit the ANOVA Model.

The following code shows how to create a fake dataset with three groups (A, B, and C) and fit a one-way ANOVA model to the data to determine if the mean values for each group are equal:

#make this example reproducible

```
set.seed(0)
```

```
#create data
```

```
data <- data.frame(group = rep(c("A", "B", "C"), each =  
30),  
values = c(runif(30, 0, 3),  
runif(30, 0, 5),  
runif(30, 1, 7)))
```

```
#view first six rows of data
```

```
head(data)
```

```
group values
```

```
1 A 2.6900916
```

```
2 A 0.7965260
```

```
3 A 1.1163717
```

```
4 A 1.7185601
```

```
5 A 2.7246234
```

```
6 A 0.6050458
```

```
#fit one-way ANOVA model
```

```
model <- aov(values~group, data=data)
```

```
#view the model output
```

```
summary(model)
```

```
Df Sum Sq Mean Sq F value Pr(>F)
group 2 98.93 49.46 30.83 7.55e-11 ***
Residuals 87 139.57 1.60
```

```
---
```

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

We can see that the overall p-value from the ANOVA table is 7.55e-11. Since this is less than .05, we have sufficient evidence to say that the mean values across each group are not equal. Thus, we can proceed to perform Tukey's Test to determine exactly which group means are different.

Step 2: Perform Tukey's Test.

The following code shows how to use the `TukeyHSD()` function to perform Tukey's Test:

```
#perform Tukey's Test
TukeyHSD(model, conf.level=.95)
```

Tukey multiple comparisons of means
95% family-wise confidence level

```
Fit: aov(formula = values ~ group, data = data)
```

```
$group
```

```
diff lwr upr p adj
```

```
B-A 0.9777414 0.1979466 1.757536 0.0100545
```

```
C-A 2.5454024 1.7656076 3.325197 0.0000000
```

```
C-B 1.5676610 0.7878662 2.347456 0.0000199
```

The p-value indicates whether or not there is a statistically significant difference between each program. We can see from the output that there is a statistically significant difference between the mean weight loss of each program at the 0.05 significance level.

In particular:


P-value for the difference in means between B and A: .0100545
 P-value for the difference in means between C and A: .0000000
 P-value for the difference in means between C and B: .0000199

We can use the `plot(TukeyHSD())` function to visualize the confidence intervals as well:

```
#plot confidence intervals
```

```
plot(TukeyHSD(model, conf.level=.95), las = 2)
```

Note: The *las* argument specifies that the tick mark labels should be perpendicular (*las=2*) to the axis.



We can see that none of the confidence intervals for the mean value between groups contain the value zero, which indicates that there is a statistically significant difference in mean loss between all three groups. This is consistent with the fact that all of the p-values from our hypothesis tests are below 0.05.

For this particular example, we can conclude the following:

The mean values of group C are significantly higher than the mean values of both group A and B. The mean values of group B are significantly higher than the mean values of group A.

A Guide to Using Post Hoc Tests with ANOVA

How to Conduct a One-Way ANOVA in R

How to Conduct a Two-Way ANOVA in R