How can I do post-hoc pairwise comparisons in R?

Authored by

stats writer

June 30, 2024

RECOMMENDED CITATION

stats writer (2024). *How can I do post-hoc pairwise comparisons in R?*. PSYCHOLOGICAL SCALES. Retrieved from https://scales.arabpsychology.com/?p=161802

Post-hoc pairwise comparisons in R refer to a statistical technique used to compare multiple groups or conditions after conducting an analysis of variance (ANOVA). This method allows for a more in-depth understanding of the differences between groups and can help identify specific factors that contribute to these differences. In R, post-hoc pairwise comparisons can be performed using various packages, such as TukeyHSD, multcomp, or emmeans, which provide functions for calculating and visualizing the results of these comparisons. This technique is commonly used in research and data analysis to further explore and interpret the findings of an ANOVA.

How can I do post-hoc pairwise comparisons in R? | R FAQ

Post-hoc pairwise comparisons are commonly performed after significant effects
have been found when there are three or more levels of a factor. After an ANOVA,
you may know that the means of your response variable differ significantly
across your factor, but you do not know which pairs of the factor levels are
significantly different from each other. At this point, you can conduct
pairwise comparisons.

We will demonstrate the how to conduct pairwise comparisons in R and the different
options for adjusting the p-values of these comparisons given the number of

tests conducted. We will be using the hsb2 dataset and looking at the variable write by ses. We will first look at the means and standard deviations by ses.

```
hsb2<-read.table("https://stats.idre.ucla.edu/stat/data/hsb2.csv", sep=",", header=T)
attach(hsb2)

ses <- factor(ses)
levels(ses) <- c("low","medium","high")
female <- factor(female)
levels(female) <- c("male","female")

tapply(write, ses, mean)
```

```
low medium high
50.61702 51.92632 55.91379
```

```
tapply(write, ses, sd)
```

```
low medium high
9.490391 9.106044 9.442874
```

One-Way ANOVA

In R, we can run the ANOVA with the aov command.

```
a1 <- aov(write ~ ses)
summary(a1)
```

```
Df Sum Sq Mean Sq F value Pr(>F)
ses 2 859 429.4 4.97 0.00784 **
Residuals 197 17020 86.4
---
Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

From this output, we can see that ses is significant in the 2-degrees of freedom test, but we do not know which pairs of ses levels are significantly different from each other. However, this will require three tests (high vs. low, high vs. middle, low vs. middle), so we wish to adjust what we consider to be statistically significant to account for this multiplicity of tests. For an one-way ANOVA (ANOVA with a single factor) We can first see the unadjusted p-values using the pairwise.t.test command and indicating no adjustment of p-values:

```
pairwise.t.test(write, ses, p.adj = "none")
```

```
Pairwise comparisons using t tests with pooled SD

data: write and ses

low medium
medium 0.4306 -
high 0.0041 0.0108

P value adjustment method: none
```

With this same command, we can adjust the p-values according to a variety of methods. Below we show Bonferroni and Holm adjustments to the p-values and others are detailed in the command help.

```
pairwise.t.test(write, ses, p.adj = "bonf")
```

```
Pairwise comparisons using t tests with pooled SD

data: write and ses
```

```
low medium
medium 1.000 -
high 0.012 0.032

P value adjustment method: bonferroni
```

```
pairwise.t.test(write, ses, p.adj = "holm")
```

```
Pairwise comparisons using t tests with pooled SD

data: write and ses

low medium
medium 0.431 -
high 0.012 0.022

P value adjustment method: holm
```

We can see that the adjustments all lead to increased p-values, but consistently the high-low and high-middle pairs appear to be significantly different at alpha = .05. The pairwise.t.test command does not offer Tukey post-hoc tests, but there are other R commands that allow for Tukey comparisons. Below, we show code for using the TukeyHSD (Tukey Honest Significant Differences).

```
TukeyHSD(a1)
```

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

Fit: aov(formula = write ~ ses)

$ses
diff lwr upr p adj
medium-low 1.309295 -2.6052575 5.223847 0.7096950
high-low 5.296772 0.9887256 9.604818 0.0114079
high-medium 3.987477 0.3296892 7.645265 0.0289035
```

We can see that these results are significant with what we saw using other adjustments for the p-values.

Two (or more) Factor ANOVA

You may be fitting an ANOVA with multiple factors. Below we look at write on ses and female.

```
a2 <- aov(write ~ ses + female)
summary(a2)
```

```
Df Sum Sq Mean Sq F value Pr(>F)
ses 2 859 429.4 5.387 0.00528 **
female 1 1398 1398.1 17.541 4.25e-05 ***
Residuals 196 15622 79.7
---
Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

We can look at pair-wise comparisons of the ses levels after adjusting for female. The TukeyHSD command still works well, though now we must specify which factor is of interest.

```
TukeyHSD(a2, "ses")
```

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

Fit: aov(formula = write ~ ses + female)

$ses
diff lwr upr p adj
medium-low 1.309295 -2.4507360 5.069325 0.6896199
high-low 5.296772 1.1587797 9.434764 0.0079527
```

```
high-medium 3.987477 0.4740753 7.500879 0.0216707
```