



EvalFest How-Tos for Conducting a One-way Analysis of Variance (ANOVA) Test using R

One-way ANOVA Test

Hi, this is Jane. In this video, we will talk about conducting an ANOVA test using R, and I will use RStudio as the editor. Let's get started.

1. Intro and Data Setup
 - a. If you've watched our video that explains the concept of a one-way ANOVA test, you may recall that a one-way ANOVA test helps you figure out when you really do have apples, oranges, pears, etc. in your data, where the different fruits are different groups of people.
 - b. In this example, we are going to take a look at average Overall Rating of the festival by Time at Festival, using four groups for Time at Festival – 31-45 minutes, 46-60 minutes, 61-90 minutes, and More than 91 minutes. The Overall Rating question asked festival attendees, "How would you rate today's event overall?" Festival goers could choose an answer on a 5-point scale from Poor to Excellent.
 - c. Here, Time at Festival is the predictor variable – we want to know if the time spent at the festival predicts differences in overall rating. Overall Rating is our outcome variable – it is the variable that we want to learn more about.
 - d. Remember that your data should be saved in either a comma separated values (.csv) or tab separated values (.tsv) format.
2. Conducting the one-way ANOVA test
 - a. Please see the file containing the code we will use to tell Rstudio to run an ANOVA test, and subsequent post hoc tests. We will talk more about post hoc tests in just a few minutes. We have made all of this code available for you so that you can download the file and then change as needed to work with your data. If you've watched our video that shows you how to set up the data and variables in R, you will notice that lines 2, 3, 6, and 7 are part of the initial setup.
 - i. Note this extra piece of code in line 3 that begins with na.strings. This code tells R to ignore any instances of missing data. This is a critical piece of code to ensure that we run the proper ANOVA and subsequent post hoc tests.
 - b. Go ahead and run lines 2, 3, 6, and 7 to set up the data for R.
 - c. Line 10 is the code that tells R to run an ANOVA test to compare Overall Rating among the 4 different Time at Festival groups.
3. Interpretation of the one-way ANOVA test
 - a. And the results of the ANOVA test appear in the Console window.
 - b. RStudio provides the degrees of freedom, the value of F, and the p-value

(labeled $\Pr(>F)$). You're most interested in the p-value to determine if you have differences between one or more pairs of Time at Festival groups. If the p-value is less than .05, that means you DO have differences. In more statistical terms, you use the p-value to determine if the mean of the outcome variable differs between at least two of your groups. For this scenario, we see that our p value is 3.52×10^{-15} . Because this is less than .05, we know that we have one or more pairs of Time at Festival groups that differ significantly in average overall rating of festival.

- c. Note that the p-value simply tells you if there are pairs of groups for which the ratings are different – it doesn't tell you which groups differ significantly from one another. To figure that out, we need to run a set of Bonferroni post hoc tests.
4. Bonferroni post hoc tests
 - a. Remember that the p-value is the risk that your results are actually from chance and not based on a real difference in your data. The value $p < .05$ is the acceptable risk for one test. If you do more than one test, you have to adjust your risk level to help make sure that your estimates account for the extra level of risk. The more tests you conduct, the greater the possibility of making the wrong decision based on the p-value. R is able to run the full set of 6 post hoc independent samples tests simultaneously, and adjusts the p-value for each post hoc test to determine whether or not the two groups being tested are significantly different from one another. Because R makes what is called the Bonferroni adjustment for you, you don't have to make this adjustment when assessing the p-value. Therefore, a p-value of less than .05 indicates that the average overall rating between the two groups being tested is different – that you DO have apples and oranges.
 - b. Line 13 contains the code that runs the 6 simultaneous independent samples t-tests, and also makes the Bonferroni adjustment to the p-values.
 - c. In the Console window, we see that those who had been at the festival more than 91 minutes rated the festival significantly differently than all other groups.
 - d. In order to see which groups rated the festival more positively, we will need to compute the average festival rating by TimeatFestival group.
 - e. Line 16 is the code that will provide the OverallRating mean for each TimeatFestival group. It is key to include `na.rm="T"`, so that any missing data will be ignored.
 - f. We can now see that those who had been at the festival more than 91 minutes rated the festival significantly *more positively* than all other groups.
 5. Final interpretation
 - a. So how do you think about interpreting these results? Remember that these ratings were made on a five-point scale, where 1 was Poor and 5 was Excellent. The average ratings were above 4.0 for all groups, so somewhere between Very Good and Excellent on the scale. When interpreting these kinds of results, you want to think about the kinds of ratings you have overall and the group differences. So, for these results, we might say something like: Ratings for the Festival were quite positive. As festival-goers spent more time at the festival, average overall rating increased. Specifically, those who had

been at the festival more than 91 minutes rated the festival significantly more positively than all other groups.



UNC
MOREHEAD PLANETARIUM
AND SCIENCE CENTER



karen peterman
CONSULTING

