R’s TukeyHSD() function

aov() function

Numeric response variable

Categorical grouping variables

```
mass.aov <- aov(antler.mass ~ diet, data = df.mass)
```

```
> TukeyHSD(mass.aov)
  Tukey multiple comparisons of means
    95% family-wise confidence level

Fit: aov(formula = antler.mass ~ diet, data = df.mass)

$diet
                 diff        lwr      upr     p adj
Hi.Lo-Hi.Hi -43.97973 -205.89517 117.9357 0.7807326
Lo.Hi-Hi.Hi 133.83308  -28.08236 295.7485 0.1198088
Lo.Hi-Hi.Lo 177.81281   15.89737 339.7282 0.0292046
```

Difference between group means
(effect size)

Lower and upper adjusted 95% confidence
intervals (CI)

Adjusted p-value

TukeyHSD stands for “Tukey’s Honestly Significant Difference”. It is related to a t-test but carries out a correction for **multiple comparisons** when all possible comparison between groups are being made. It is recommended that this technique (or one like it) be used when carrying out exploratory “unplanned comparisons” that were not specified prior to designing the experiment or gathering the data (“post hoc comparisons”).

Unlike the pairwise.t.test() function, the TukeyHSD() function takes just a single argument, which is the output of the function aov(). aov() stands for analysis of variance and is a way to fit an ANOVA model R.

TukeyHSD() produces much more output than pairwise.t.test(). In particular, it gives you the **difference between the means** of each group (“diff”) and a **95% confidence interval** (CI) around this difference that has been adjusted for multiple comparisons. The more comparisons, the winder the interval becomes, and the large the adjusted p-value (adj) gets)

95% CIs that contain zero will have a p value > 0.05. The hypothesis that the difference between two means is 0 therefore cannot be rejected

Data used in example

```
 summary(df.mass)
 antler.mass         diet
Min.   :346.0   Hi.Hi:10
1st Qu.:526.9   Hi.Lo:10
Median :648.4   Lo.Hi:10
Mean   :635.9
3rd Qu.:762.4
Max.   :919.3
```

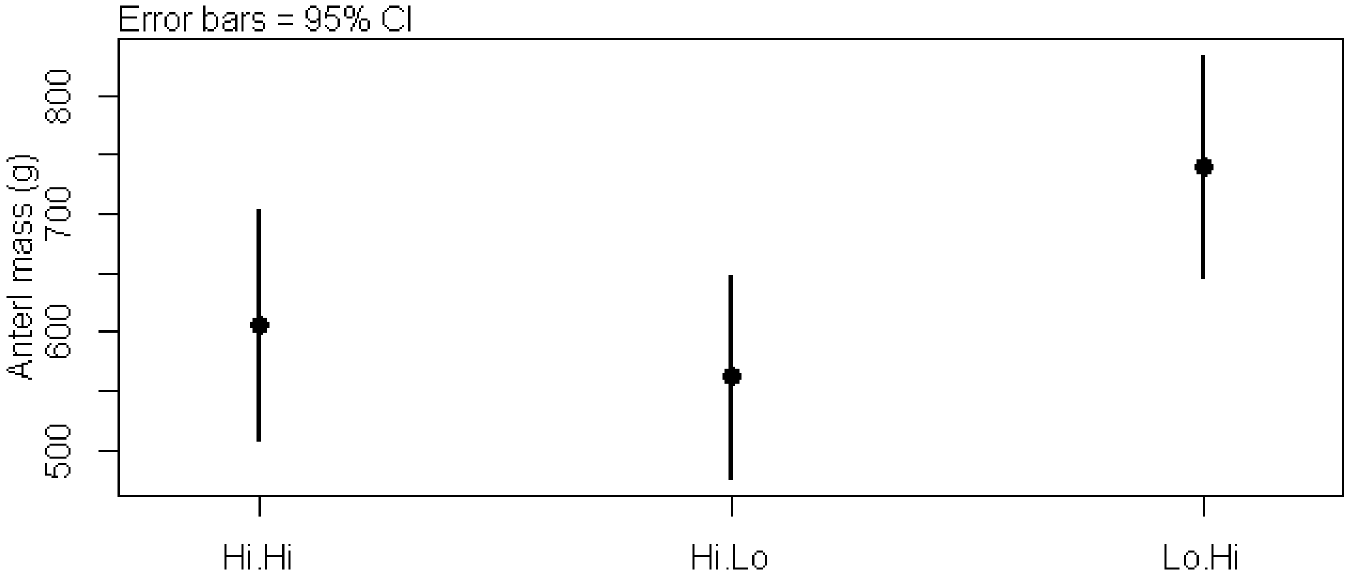

The data used in this example is a dataframe called “df.mass” that has 2 columns. The data are from an experiment on how different periods of high and low protein affect antler growth in white-tailed deer. A numeric variable is in the 1st column “antler.mass”, and a categorical variables (“diet”) with 3 groups is in the 2nd column. “Hi.Hi” is a diet with consistently high protein content. The “Hi.Lo” treatment provided high protein followed by a period with a low-protein diet. “Lo.Hi” began with a low protein diet and switched to a high protein diet. The plot to the right shows the raw means.