

Assignment 2—Introduction to Statistics Using R

Paul Gribble
Winter, 2017

* Due Sunday February 5

1 Single Factor Between Subjects ANOVA

The .csv file called `migraine.csv` contains 27 pain ratings collected from migraine sufferers following the ingestion of one of three drugs after the onset of migraine (drug A, B or C). Pain ratings are on a scale of 1 (least pain) to 10 (most pain).

Load the data into R with the following code:

```
fname <- "http://www.gribblelab.org/stats/data/migraine.csv"
migraine <- read.table(fname, sep=";", header=TRUE)
```

1. Generate some kind of graphical display of the data. It's up to you what kind. Label the axes and give the Figure a title.
2. Perform a one-way analysis of variance to determine if there is evidence that pain is affected by drug. You can use the following commands in R if you like:

```
results <- aov(pain ~ drug, data=migraine)
summary(results)
```

3. Where appropriate, test the assumptions of ANOVA. You can use `bartlett.test` to test the homogeneity of variance assumption. You can use the `shapiro.test` to test the normality assumption. You can do the test either on all the data lumped together, or on each group separately—your choice.
4. With reference to the appropriate elements of the ANOVA output, what is your interpretation of the analysis?
5. The omnibus test is significant ($F(2, 24)=11.91$, $p<0.000256$). Now test which groups are different from each other. Assume you haven't planned comparisons in advance. Conduct all possible pairwise tests, and be sure to apply an adjustment for Type-I error. Explain

why you chose your particular adjustment.

Hint: the `TukeyHSD` function in R can be applied to the output of the `aov` function (in the above code, the variable called `results`).

6. On the basis of your analyses? What can you conclude about each drug's affect on migraine pain?