

Hacker Dojo Machine Learning

Homework 1

Mike Bowles, PhD & Patricia Hoffman, PhD.

(attribution to Professor David Mease)

1) This question uses the data at http://www.cob.sjsu.edu/mease_d/bus297D/myfirstdata.csv. Download it to your computer.

a) Read in the data in R using `data<-read.csv("myfirstdata.csv",header=FALSE)`. Note, you first need to specify your working directory using the `setwd()` command. Determine whether each of the two attributes (columns) is treated as qualitative (categorical) or quantitative (numeric) using R. Explain how you can tell using R.

b) What is the specific problem that causes one of these two attributes to be read in as qualitative (categorical) when it seems it should be quantitative (numeric)?

c) Use the command `plot()` in R to make a plot for each column by entering `plot(data[,1])` and `plot(data[,2])`. Because one variable is read in as quantitative (numeric) and the other as qualitative (categorical) these two plots are showing completely different things by default. Explain exactly what is being plotted in each of the two cases. Include these two plots in your homework.

d) Read the data into Excel. Excel should have no problem opening the file directly since it is .csv. Create a new column that is equal to the second column plus 10. What is the result for the problem observations (rows) you identified in part b? What specific outcome does Excel display?

2) This question uses the data at http://www.cob.sjsu.edu/mease_d/bus297D/twomillion.csv. Download it to your computer.

a) Read the data into R using `data<-read.csv("twomillion.csv",header=FALSE)`. Note, you first need to specify your working directory using the `setwd()` command. Extract a simple random

sample with replacement of 10,000 observations (rows). Show your R commands for doing this.

b) For your sample, use the functions `mean()`, `max()`, `var()` and `quantile(.,.25)` to compute the mean, maximum, variance and 1st quartile respectively. Show your R code and the resulting values.

c) Compute the same quantities in part b on the entire data set and show your answers. How much do they differ from your answers in part b?

d) Save your sample from R to a csv file using the command `write.csv()`. Then open this file with Excel and compute the mean, maximum, variance and 1st quartile. Provide the values and name the Excel functions you used to compute these.

e) Exactly what happens if you try to open the full data set with Excel?

3) This question uses a sample of 1500 California house prices at http://www-stat.wharton.upenn.edu/~dmease/CA_house_prices.csv and a sample of 10,000 Ohio house prices at http://www-stat.wharton.upenn.edu/~dmease/OH_house_prices.csv. Download both data sets to your computer. Note that the house prices are in thousands of dollars.

a) Use R to produce a single graph displaying a boxplot for each set (hint: `boxplot(data[,2],data[,3],col="blue",`

`main=" House Boxplot ",`
`names=c("CA houses","Ohio houses"),ylab="Prices")`

). Include the R commands and the plot. Put a name in the title of the plot (for example, `main="House Boxplots"`).

b) Use R to produce a frequency histogram for only the California house prices. Use intervals of width \$500,000 beginning at 0 and ending at \$3.5 million. Include the R commands and the plot. Put your name in the title of the plot.

c) Use R to plot the ECDF of the California houses and Ohio houses on the same graph (see example2.r). Include a legend. Include the R commands and the plot. Put your name in the title of the plot.

4) This question uses the data at <http://www-stat.wharton.upenn.edu/~dmease/football.csv>. Download it to your computer. This data set gives the total number of wins for each of the 117 Division 1A

college football teams for the 2003 and 2004 seasons.

a) Use `plot()` in R to make a scatter plot for this data with 2003 wins on the x-axis and 2004 wins on the y-axis. Use the range 0 to 12 for both the x-axis and y-axis. Include the R commands and the plot. Put your name in the title of the plot.

b) Why are there fewer than 117 points visible on your graph in part a? Describe the solution we discussed in class to deal with this problem (but don't actually do it).

c) Compute the correlation in R using the function `cor()`.

d) How does the value in part c change if you add 10 to all the values for 2004?

e) How does the value in part c change if you multiply all the 2004 values by 2?

f) How does the value in part c change if you multiply all the 2004 values by -2?

5) This question uses the sample of 10,000 Ohio house prices at http://www-stat.wharton.upenn.edu/~dmease/OH_house_prices.csv. Download the data set to your computer. Note that the house prices are in thousands of dollars.

a) What is the median value? Is it larger or smaller than the mean?

b) What does your answer to part a suggest about the shape of the distribution (right-skewed or left-skewed)?

c) How does the median change if you add 10 (thousand dollars) to all the values?

d) How does the median change if you multiply all the values by 2?

5) This question uses the following people's ages: 19,23,30,30,45,25,24,20. Store them in R using the syntax `ages<-c(19,23,30,30,45,25,24,20)`.

a) Compute the standard deviation in R using the `sd()` function.

- b) Compute the same value by hand and show all the steps.
- c) Using R, how does the value in part a change if you add 10 to all the values?
- d) Using R, how does the value in part a change if you multiply all the values by 100?