

BIOS: 4120 Lab 5

February 13-14, 2018

Objectives

In today's lab we will:

1. Review & expand what we've covered in R
2. Learn how to construct a Scatter plot
3. Explore linear regression and correlation functions

R Practice and Review:

Here are a couple of questions that you should be able to answer based on what we have learned so far in lab.

1. How would you read in the Tips dataset from the course website?

```
tips<-read.delim("http://myweb.uiowa.edu/pbreheny/data/tips.txt")
```

2. How do you access just the Tip column? Can you find the mean?
3. How do you access just the Tip column but only for the smokers? What's the mean tip from smokers?
4. How do you access just the Tip column but only when the total bill is less than 15? What is the mean for this group?

The answers regarding the means are presented below.

```
mean(tips$Tip)
```

```
## [1] 2.998279
```

```
mean(tips$Tip[tips$Smoker=="Yes"])
```

```
## [1] 3.00871
```

```
mean(tips$Tip[tips$TotBill<15])
```

```
## [1] 2.05025
```

Manipulating Columns

You can perform operations on whole columns of data at a time. For instance:

```
tipPctgs<-tips$Tip/tips$TotBill  
summary(tipPctgs)
```

```
##   Min. 1st Qu.  Median    Mean 3rd Qu.    Max.  
## 0.03564 0.12913 0.15477 0.16080 0.19148 0.71034
```


```
tipPctgs2<-tipPctgs*100  
summary(tipPctgs2)
```

```
##   Min. 1st Qu.  Median    Mean 3rd Qu.    Max.  
##   3.564  12.913  15.477  16.080  19.148  71.034
```

Scatter Plots

When we are interested in visualizing the connection between two continuous variables we can use a scatter plot which is done using the following function:

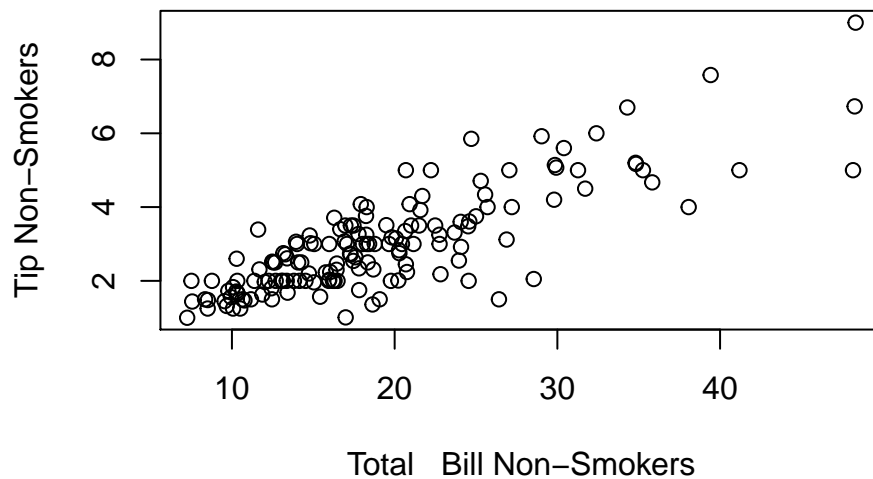
```
plot(tips$TotBill, tips$Tip)
```



Looking at this plot we can see a positive association between the variables and that there is a good amount of variation. Also, you can see horizontal lines are beginning to form. *What's causing these lines? Why would the scatter plot not be used to compare other variables in this dataset?*

Maybe we're only curious about the total bill compared to the tip for individuals who don't smoke. Here's how we could plot that:

```
plot(tips$TotBill[tips$Smoker=="No"], tips$Tip[tips$Smoker=="No"],
     xlab = 'Total Bill Non-Smokers', ylab = 'Tip Non-Smokers')
```



Notice that we were able to make reading the plot easier by changing the x and y axis labels using the parameters `xlab` and `ylab`.

Constructing Models

In R, the function we will use is called ‘`lm`’, which stands for “linear model”. Regression is useful because it can be generalized to all kinds of settings through the notion of a model, as alluded to by its name in R.

To create a model in R, the code looks like this:

```
mod <- lm(Tip~TotBill,data=tips)
(mod)
```

```
##
## Call:
## lm(formula = Tip ~ TotBill, data = tips)
##
## Coefficients:
## (Intercept)      TotBill
##      0.9203         0.1050
```

Note: The general function should look like `lm(y variable ~ x variable, data = the dataset)`.

Printing out the model itself gives you the intercept and slope. This specific output tells us that for every additional dollar that a meal costs, the waiter can expect to get 10.5 cents more on his tip.

We could have also calculated this using the output from `cor()` function. By itself, the `cor()` function calculates the correlation coefficient between two vectors (shown below using the tip and total bill amounts). When we multiply that output by the ratio of the standard deviations of those two variables, we end up with the slope of the model.

```
cor(tips$Tip, tips$TotBill)
```

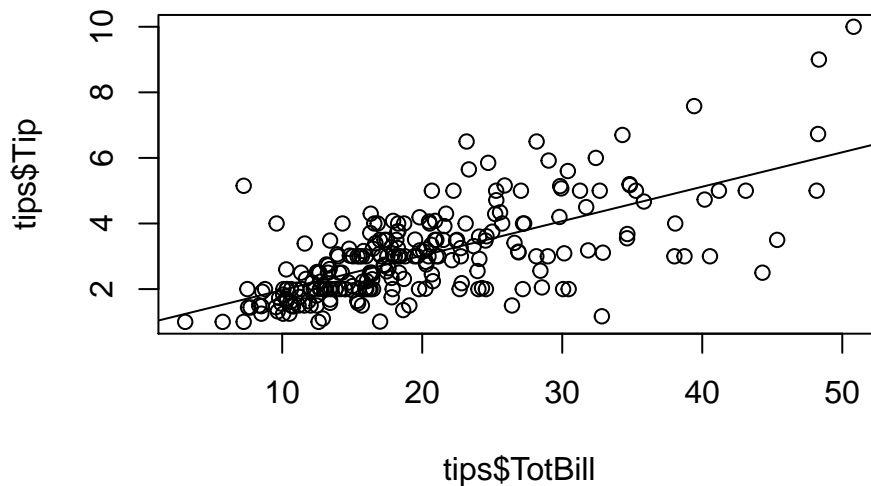
```
## [1] 0.6757341
```

```
cor(tips$Tip,tips$TotBill)*sd(tips$Tip)/sd(tips$TotBill)
```

```
## [1] 0.1050245
```

If you are interested in adding this regression line to a plot of the data, you can use the `abline()` function, and just put the name of your model inside the parentheses.

```
plot(tips$TotBill, tips$Tip)  
abline(mod)
```



Example:

- Make a scatter plot of tip amount vs total bill for only individuals who smoke.
- Now add a regression line to the plot.

Now let's shift back to looking at the entire dataset.

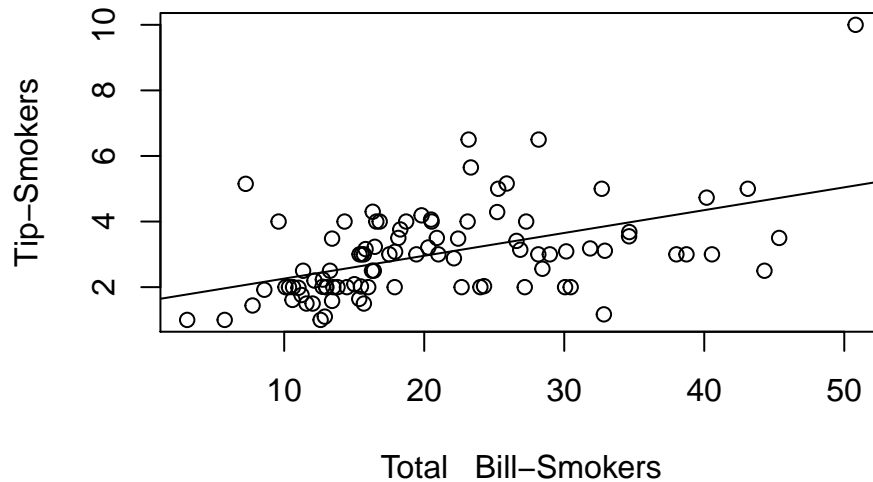
- Suppose a table is 1 standard deviation above average in terms of total bill. How many standard deviations above average in terms of tip would you expect it to be?
- Suppose a table is \$2 above average tip. How many dollars above the average total bill would you expect it to be?
- Suppose a table is \$10 above the average total bill. What would we expect the tip to be?

```

# Part A
plot(tips$TotBill[tips$Smoker=="Yes"], tips$Tip[tips$Smoker=="Yes"],
     xlab = 'Total Bill-Smokers', ylab = 'Tip-Smokers')

# Part B
smokermodel <- lm(Tip[tips$Smoker=="Yes"]~TotBill[tips$Smoker=="Yes"],data=tips)
abline(smokermodel)

```



```

# Part C
cor(tips$TotBill,tips$Tip)

```

```
## [1] 0.6757341
```

```

# Part D
Zx <- 2/sd(tips$Tip)
Zy <- Zx * cor(tips$Tip,tips$TotBill)
Zy * sd(tips$TotBill)

```

```
## [1] 8.695428
```

```

# Part E
Zx <- 10/sd(tips$TotBill)
Zy <- Zx * cor(tips$TotBill,tips$Tip)
(y <- mean(tips$Tip) + Zy * sd(tips$Tip))

```

```
## [1] 4.048524
```

```

# Alternatively, we can use the values from the 'mod' model we created earlier
mod$coefficients[1] + mod$coefficients[2]*(mean(tips$TotBill)+10)

```

```

## (Intercept)
## 4.048524

```