

STATS 60 Summer 2020: HW3

Remarks

Please show your work, and briefly explain your answers.

If a problem has multiple parts, each will be graded separately. We will not subtract grades if your solution to one part is correct but the number you use from previous parts is wrong.

You will need to use the packages `ggplot2`, `dplyr` and `lubridate` for this problem set. Install the package `lubridate` with `install.packages("lubridate")` and load the package with `library(lubridate)`

```
library(ggplot2)
library(dplyr)
library(lubridate)
```

Problem 1 Reading habits

Your friend wants find out reading habits of residents of Palo Alto. In particular, how many books people read per year and which types of books they read most.

- (a) He has access to the checkout records at Palo Alto library, which contain the following variables

Year Which year is it?

Genre Book category, such as “Action”, “Anthology”, “History” etc.

Total Number of checkouts

Describe the data type and measurement scale of each of these variable. Is it qualitative or quantitative? Is it nominal, ordinal, interval or ratio?

- (b) He divides the total checkouts by the total visitors to the library in each year and obtain the ratio. He also finds the most popular categories in terms of this ratio. Do you think his approach answers the questions? Explain your answers in a paragraph. If yes, explain why. If the answer is no, explain what more information you need to get a better answer.

Problem 2 Pooled testing

Recently there’s some discussions about group screening of COVID-19. By pooling samples from many people into groups, and evaluating pools rather than individuals, we can potentially reduce the number of tests and increase the number of people tested. This problem explores the cost-effectiveness of this idea.

Suppose a test has sensitivity 99.9% and specificity 99.8%. For simplicity, we assume that this is true for pooled tests as well, that is to say, we can treat the pooled sample as a single sample, which is infected if at least one person in the group is infected. In this problem, we consider groups of size 10. Finally assume the infection probability is 14 in 1000 people.

- (a) What’s the probability that everyone is healthy in a group?
- (b) What’s the probability a group tests positive?
- (c) Now suppose a group tests positive, what’s the posterior odds that at least one person from the group is infected?

- (d) Suppose we test 5000 separate groups, which corresponds to testing 50,000 people. What's the probability that 650 groups test positive? What about 800?

Extra How many groups do you expect to test positive?

Problem 3 Sampling distribution of the mean

The Current Population Survey at this link provides summary data of individual education, work and income. Specifically, we use the household income data from year 2018 at this link, which contains the number of households (in thousands) in each income bracket. For example, if you download the data table, you can see there are 4,283 (in thousands) households out of 128,581 whose yearly income is less than \$5000.

From this source, we generate a hypothetical population of size 128,581, i.e. we take one per thousand, assuming the income of all the households in one bracket is the midpoint of that bracket. Finally, for households in the last bracket "\$250,000 and above", we treat their income as 250,000. The following R code creates this population and store it in the vector `all_income`, whose length is 128,581 and contains hypothetical incomes. We will work with this hypothetical income population `all_income`.

```
income <- c(seq(0, 200000, by = 5000), 250000) # income brackets
mid_income <- c((income[-42] + income[-1]) / 2, 250000) # mid-points of the bracket
N <- c(4283, 3337, 5510, 5772, 5672, 5469, 5822, 5404, 5195, 4839,
      5300, 4417, 4604, 3999, 3795, 3950, 3349, 3064, 3102, 2581,
      2866, 2449, 2318, 1971, 2004, 1780, 1678, 1426, 1414, 1316,
      1492, 978, 1161, 970, 905, 835, 772, 686, 584, 565,
      4572, 6375) # number of households in each bracket
all_income <- rep(mid_income, N)
```

- (a) Make a histogram of this hypothetical income population. Why is there a uptick at 250,000? Is it from a normal distribution? Why?
- (b) What is the population mean, median and interquartile range?
- (c) What is the z-score of a household income of \$40,000?
- (d) Draw a random sample of size 1000 of all the incomes, what is the mean of this random sample? In R, we can draw a random sample using the `sample` function, for example the following R code draws a random sample without replacement from the population and stores it to the vector `s`.

```
s <- sample(all_income, size = 1000)
```

- (e) Draw 1000 random samples of size $n = 1000$, compute the mean of each sample. Draw a histogram of all the sample means. Is it approximately normal? What is the mean and standard deviation of these sample means?

You can use a for loop iterate a process many times. Here's an example.

```
result <- numeric(length = 10) # a vector of 100 zero
for(i in 1:10){ # iterate i from 1 to 100
  # write the action to perform in one loop inside of bracket { }
  # here we assign the i-th element of result to i
  result[i] <- i
}
result
```

```
## [1] 1 2 3 4 5 6 7 8 9 10
```

(f) Repeat part (d) and (e) with sample size $n = 5000$, what do you observe?

Problem 4 Drought data

This exercise explores more the US drought data we looked at in the Rlab. The data we use was downloaded from US drought monitor website, the state drought level in terms of percent area from 2015-01-01 to 2019-12-31. You can download the dataset in the course website assignment webpage and read the data into a data frame `drought` using the following code. Make sure you change the path to where your data is located.

```
drought <- read.csv("~/Desktop/drought.csv") # use this path if your data is at "Desktop" folder
```

Before we start, we modify column names and select columns of interest. For simplicity, we assume each row to correspond to 7 day drought level starting at `ValidStart`. We convert `ValidStart` to a `Date` object and store it at the column `date`.

```
drought <- drought %>%  
  mutate(  
    date = ymd(ValidStart), # ymd is a function from lubridate package  
    state = as.character(StateAbbreviation)  
  ) %>%  
  select(date, state, None, D0:D4)
```

You can take a quick look at the data with `glimpse`.

```
glimpse(drought)
```

```
## Observations: 13,624  
## Variables: 8  
## $ date <date> 2019-12-31, 2019-12-24, 2019-12-17, 2019-12-10, 2019-12-03, ...  
## $ state <chr> "AK", "AK", "AK", "AK", "AK", "AK", "AK", "AK", "AK", "AK", "...  
## $ None <dbl> 93.18, 93.18, 93.18, 92.48, 92.48, 91.17, 91.17, 90.29, 90.29...  
## $ D0 <dbl> 6.82, 6.82, 6.82, 7.52, 7.52, 8.83, 8.83, 9.71, 9.71, 9.71, 1...  
## $ D1 <dbl> 0.83, 0.83, 0.83, 0.83, 0.83, 0.83, 0.83, 3.96, 3.96, 3.96, 4...  
## $ D2 <dbl> 0.00, 0.00, 0.00, 0.00, 0.00, 0.00, 0.00, 0.00, 0.00, 2.00, 2...  
## $ D3 <dbl> 0.00, 0.00, 0.00, 0.00, 0.00, 0.00, 0.00, 0.00, 0.00, 0.00, 0...  
## $ D4 <dbl> 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0...
```

(a) Drought severity in each state

In this exercise we will compute the average drought severity in each state in the year 2019. The drought severity index at this link is given by

$$1(D_0) + 2(D_1) + 3(D_2) + 4(D_3) + 5(D_4) = DSCI$$

- (1) Filter the data to California.
- (2) Filter the California data to the year 2019. To extract year, you can use

```
year(drought$date)
```

- (3) Compute the average DSCI of California in the year 2019.
- (4) Compute the average DSCI of every state in the year 2019.
- (5) Display your findings in part (4) in a visualization.

(b) Is drought seasonal?

- (1) In the state of California, compute the average drought severity index (see question (a) part (1)) for each month in the year 2019. Which month sees the most severe drought? Use `month()` to extract month from a `Date` object.

```
month(drought$date)
```

- (2) Is your finding in part (1) consistent in all five years from 2015 - 2019?

(c) Your question

Describe a question you have about drought levels and use summary statistics or a visualization from this data to answer it.