

## Assignment 2

Due: October 13, 2017 @ 11:59pm

### Instructions

In the [Github Classroom repository](#) that you copied, you should have a template file named `assignment_2.Rmd`. Use this file to answer the [questions below](#). Also fill out any missing information at the top of the RMarkdown document. Once you're done, be sure to save, commit, and push to Github. Then, [submit a pull request](#).

### Overview




Figure 1: A photograph of the *Titanic* leaving Southampton on April 10, 1912.

For this homework assignment, you will be exploring a dataset about the passengers on the *Titanic*, the British passenger liner that crashed into an iceberg during its maiden voyage and sank early in the morning on April 15, 1912. The tragedy stands out as one of the deadliest commercial maritime disasters during peacetime in history. More than half of the passengers and crew died, due in large part to poor safety standards, such as not having enough lifeboats or not ensuring all lifeboats were filled to capacity during evacuation.

This dataset presents the most up-to-date knowledge about the passengers that were on the *Titanic*, including whether or not they survived. This dataset is frequently used to introduce using machine learning techniques that take multiple inputs and use them to predict an outcome, in this case whether a passenger is likely to have survived. While we won't be using a machine learning model in this assignment, there is still a lot of information that can be learned by exploring the dataset using the `tidyverse` suite.

The dataset is included in your [Github starter code repository](#) for the assignment.

### About the dataset

The following are the variable (column) descriptions for the dataset:<sup>1</sup>

Variable	Description
<code>pclass</code>	Passenger Class (1 = 1st; 2 = 2nd; 3 = 3rd)
<code>survival</code>	Survival (0 = No; 1 = Yes)
<code>name</code>	Name
<code>sex</code>	Sex
<code>age</code>	Age
<code>sibsp</code>	Number of Siblings/Spouses Aboard
<code>parch</code>	Number of Parents/Children Aboard

<sup>1</sup><http://biostat.mc.vanderbilt.edu/wiki/pub/Main/DataSets/titanic.html>

Variable	Description
ticket	Ticket Number
fare	Passenger Fare (British pound)
cabin	Cabin
embarked	Port of Embarkation (C = Cherbourg; Q = Queenstown; S = Southampton)
boat	Lifeboat
body	Body Identification Number
home.dest	Home/Destination

Also note that the following definitions were used for `sibsp` and `parch`:

Label	Definition
Sibling	Brother, Sister, Stepbrother, or Stepsister of Passenger Aboard Titanic
Spouse	Husband or Wife of Passenger Aboard Titanic (Mistresses and Fiances Ignored)
Parent	Mother or Father of Passenger Aboard Titanic
Child	Son, Daughter, Stepson, or Stepdaughter of Passenger Aboard Titanic

## Questions

- When reading in the dataset using `read_csv(file = "titanic_dataset.csv")`, several of the columns are converted into inconvenient data types. Fix this so that your later analysis does not run into problems. Use the `col_types = cols()` argument within `read_csv()`, [see this section of R for Data Science for a review](#), to change the data type defaults for the following columns:
  - Convert `survived` to the logical data type
  - Convert `pclass` to the character data type
  - Convert `sibsp` to the character data type
  - Convert `parch` to the character data type
- Compute how many known passengers were on the Titanic. *Do not just print the table, use a function to count the passengers.*
- A famous directive for evacuating the *Titanic* was “women and children first”. Use your `dplyr` functions to verify the first part of this statement by counting the number of men and women that survived and that died. Then, using those counts, calculate the fraction of women that survived,

$$\frac{\text{Number of female survivors}}{\text{Total number of female passengers}}$$

and the fraction of men that survived,

$$\frac{\text{Number of male survivors}}{\text{Total number of male passengers}}$$

Do your computations support the idea that women were more likely to survive? Why or why not?

- Verify the second part of the “women and children first” directive. This will not be as straightforward as it was in the previous question, as the dataset only contains people’s ages, which can take on many values. By default, there are no columns with labels of *child* or *adult*, so you will need to create your own.

Create a new column named `child_or_adult` that uses the age data to label each passenger. For our purposes, we want to label anyone aged 0–9 as a child and anyone age 10 and up as adults. If the age cell is blank (NA) for a passenger, also label them as an “adult”. Assign this updated dataset to the variable `titanic_age_groups`.

**Hint:** You will need to use the `ifelse()` function to complete this task. An example usage of `ifelse()` is the following:

```
titanic %>%
  mutate(cheap_or_expensive = ifelse(test = fare < 15,
                                     yes = "cheap ticket",
                                     no = "not cheap"))
```

To handle blank entries, you will also need to use `is.na()` somewhere inside your `ifelse()` test.

5. Using the `titanic_age_groups` dataset you created in the previous question, count the number of children that survived and the number that did not. Do your computations support the idea that children were also more likely to survive? Why or why not?
6. A passenger’s age group and sex are not the only predictors of survival. For example, social standing and wealth can play a factor in survival. One of the parameters within this dataset acts as a proxy for distinguishing between the upper and lower classes. Which parameter is it? How do you know?
7. Group your dataset by sex and the variable you determined in question 6 and count the number that survived and the number that did not. Create a bar chart that summarizes the data, where `survived` is along the horizontal axis and the passenger counts are along the vertical axis. Use the bar chart `fill = aesthetic` to break the bar charts down by your variable from question 6. Additionally, facet over the sex variable. Interpret this visualization and describe any survival patterns that you notice.
8. Create two visualizations:
  - The first visualization should be a bar chart displaying the fraction of the passengers that survived for different values of `parch`,

$$\frac{\text{For a given parch, the number of survivors}}{\text{Total number of passengers}}.$$

Doing this requires grouping your data properly, counting the number of passengers in each grouping, and then dividing this by the total number of passengers on the ship.

- The second visualization should be a bar chart displaying the fraction of the passengers that survived for different values of `sibsp`,

$$\frac{\text{For a given sibsp, the number of survivors}}{\text{Total number of passengers}}.$$

Like above, doing this requires grouping your data properly, counting the number of passengers in each grouping, and then dividing this by the total number of passengers on the ship.

Interpret the patterns that you see in the visualizations.

9. Based on your analysis, write a list of the factors that affected the chances of survival for each passenger. You should be able to identify 4 different attributes that had a noticeable impact on survival. Justify each attribute that you list by referencing back to a table or visualization you created in a previous question.