

Instructions

Upload a PDF file, named with your UC Davis email ID and homework number (e.g., sfrei_hw3.pdf), to Gradescope (accessible through Canvas). You will give the commands to answer each question in its own code block, which will also produce output that will be automatically embedded in the output file. All code used to answer the question must be supplied, as well as written statements where appropriate.

All code used to produce your results must be shown in your PDF file (e.g., do not use `echo = FALSE` or `include = FALSE` as options anywhere). Rmd files do not need to be submitted, but may be requested by the TA and must be available when the assignment is submitted.

Students may choose to collaborate with each other on the homework, but must clearly indicate with whom they collaborated.

Gapminder

We will be using the `gapminder` dataset; you can load this by installing the package `gapminder` and then loading it as a library. Inspect the tibble by typing `?gapminder` in the console.

```
library(gapminder)
```

In this homework we aim to plot the minimum, median, and maximum life expectancy, population, and gdp per capita per continent over time. We need to first do a bit of data cleaning and preparation, and then we can use the power of ggplot.

Part (a): 5 points

Write a function `summary_stats` which takes in a tibble and a variable name and returns a tibble with the following summary statistics for the variable:

- minimum
- maximum
- median

Allow for an argument, `na.rm=`, which will specify whether or not the computation of these summary statistics will remove NA's or return NA's for any NA's in the computation. The function must use the `summarize()` function with `.groups = "drop"`.

```
# code here
```

If your code is correct, running `summary_stats(flights, dep_time)` should return

min	median	max
1	1401	2400

Part (b): 5 points

Using `across()`, compute the minimum, median, and maximum for each of the columns `lifeExp`, `pop`, and `gdpPercap` per year and per continent (i.e., min/median/max of each of these variables per year and per continent - you are doing these operations over different countries in every year-continent pair.). Save the resulting tibble as `gapminder_summary`, and print the first few rows of the tibble by writing `gapminder_summary`.

```
# code here
```

Part (c): 5 points

Make the `gapminder_summary` table in “long” format so that we have variable names `continent`, `year`, `lifeExp`, `pop`, and `gdpPercap`, and `measurement`, where `measurement` is either “min”, “max”, or “median”. Call the tibble `gapminder_summary_long`, and print the first few rows of the tibble.

```
# code here
```

Part (d): 15 points

Create three ggplots, one for each of life expectancy, population, and gdp per capita. In each plot, we want to have 5 subplots (using facet wrap or facet grid), one for each of the continents. In each subplot, we want three lines: the minimum of the variable (either `lifeExp` / `pop` / `gdpPercap`), the median of the variable, and the maximum of the variable. Do this by creating a function `plot_min_med_max()`, which takes as its input the variable name, y-axis label, and a plot title and returns a ggplot which has the 5 subplots in it. Your final 3 plots can be created by calling this function three times with inputs `lifeExp`, `pop`, and `gdpPercap` (no quotes!) and with appropriate y-axis label and plot title. Ensure that the plots have the following properties:

- The x-axis should be the same across subplots in a given plot, but the y-axis should scale separately for each sub-plot so that it is easy to visualize what is happening within each continent
- The y-axis should be a human readable form: not “`lifeExp`”, “`pop`”, or “`gdpPercap`”, but something like “Life Expetancy (in years)”, “Population”, etc.
- There is a title to the plot which describes what the subplots describe at a high-level
- The minimum line should be in red, the median in black, and the maximum in green.
- The function `plot_min_med_max()` takes 3 arguments: variable name (not a string, just a sequence of characters), y axis label (a string), and a title (a string).
- All x-axis and y-axis labels are human-readable; you may need to adjust font sizes or the orientation of the labels (look up the `theme()` function, and inspect the `axis.text.x` argument and use the function `element_text()` to do this.

```
# code here
```

Part (e): 5 points

Comment on any unusual trends that you notice in the life expectancy across continents. Which (country, year) pairs are responsible for any apparent outliers? Do you have any hypotheses for what might have caused them? Provide any R code which you use to do this.

```
# code here
```