

Stats 701 Assignment #2

Josh Errickson

Due April 7, 2017, submitted via Canvas.

Libraries I used in my solutions:

```
library(stringr)
library(jsonlite)
library(ggplot2)
library(rvest)
library(tidyr)
library(mice)
library(testthat)
```

Instructions:

- Submit **both** your RMarkdown file **and** your output file (pdf or html, your choice). The Rmarkdown should compile and show all work.
- If you use an online resource or collaborate with classmates, provide attribution.
- If you cannot finish a question, show your work and explain/demonstrate what is causing your issue for partial credit.
- Compared to assignment 1, my solutions for these exercises are substantially longer, though both are under 100 lines. Again, only provided for a reference; if you're going far beyond that without seeing the end, come see me to clarify.

Exercise 1

On this github page, <https://github.com/jhallen/joes-sandbox/tree/master/editor-perf>, user jhallen ran some performance tests on various text editors, producing a number of tables. There are two “sets”: the first measures how much memory (RSS) loading files with or without syntax highlighting takes; the second times performing an operation. However, it's hard to get an overview because the results are not combined (e.g. nano loads a file and jumps to the end very quickly, but is one of the slowest to perform a search and replace).

Using ggplot2, produce some graphics which combine these results to tell a fuller story across all browsers. You can graph the RSS and time results separately (though if you can come up with a clever way of combining them that doesn't produce visual overload, I'll give some extra credit for that).

Extra Credit: Some of the tests are not performed on some editors or doesn't run to completion (e.g look at the “Time to load 3 GB file” table). Include that information in your visualization.

Notes:

- There is no correct answer here.
- I was able to scrape this with rvest. However, it may be faster for you to scrape this manually with copy/paste into text or Excel. This is fine, but make sure you describe in words what you did (since there's no corresponding R code).
- It may be useful (but definitely not necessary) to work on the log scale (or $\log(x + 1)$ to ensure all values are positive).
- Ignore all comments (drop the Comments column, make comments in variables NA) with the exception of the few “at least X minutes” - Convert those to their max (e.g. “at least 5 minutes” -> 300 seconds).

Tips:

- The `stat` and `position` options may be useful.
- The `gather` function from the “tidyr” package may be useful for converting from wide to long data (A lot of geom’s that I tested required long data). The basic syntax is:

```
gather(<data>, <key colname>, <value colname>, <variables>)
```

You make up the `<key colname>` and `<value colname>`, and the list of `<variables>` is the variables to make long (all other variables are constant) and is just a vector of *unquoted* variable names.

Exercise 2

Grab data on all automotive facilities in Southeast Michigan from here: http://maps-semcog.opendata.arcgis.com/datasets/e71268f136dc4538abb4cffb2594277c_0 (Download Dataset on the right below the map.) This contains data on all 51 automotive plants in the area. Let’s focus on EMPLOYMENT, YEARBUILT, SQUAREFEET and attempt to answer the question: Are more recent plants built more efficiently to accommodate more workers, controlling for the size of the facility.

There are probably better statistical approaches to determine this, but for our purposes, we will address this with a linear model predicting EMPLOYMENT based upon YEARBUILT and SQUAREFEET. (Take `log` transforms of both EMPLOYMENT and SQUAREFEET.)

```
lm(log(EMPLOYMENT) ~ log(SQUAREFEET) + YEARBUILT, data = ...)
```

- 1) Clean up this data. Restrict yourself to currently operating plants. Take special note of the COMMENTS.
- 2) Do you think the missing data is MCAR, MAR or MNAR? Remember, we can only answer this by considering the data collection methods.
- 3) Use “mice” to multiply impute these three columns. Hint: Be sure to clean the data first!
- 4) Report the coefficients from $2+k$ regression models:
 - a) The original data.
 - b) Each of your k imputations.
 - c) The pooled model.

You only need do $k = 5$ iterations. Feel free to do more if you’d like, but I don’t want a massive table!

Extra Credit: Create some sort of visualization for these coefficients, specifically visualizing the original and pooled estimates relative to the distribution of imputation estimates. Use $k \geq 100$ if attempting this. There are certainly many ways to accomplish this; my solution involved `facet_wrap`.

Extra Credit 2: Create a scatter plot between `log(SQUAREFEET)` and `log(EMPLOYMENT)`, plotting all the least square lines. Clearly distinguish the fit from the original data and the pooled results versus the various imputations. This is a *much* trickier extra credit that probably isn’t worth the time! (At least, it took me nearly an hour to solve this.)

Notes:

- When cleaning small data, its often quicker to manually make modifications, either in Excel or by directly referencing the cells, e.g. `data[4,6] <- ...`. Be sure to document this in your submission!
- Take special note of COMMENTS!
- If you make decisions how to handle bad data, be explicit about what decision you make.
- Different variables can have different types of missingness.
- I may not have said this during the MI lecture, but there are arguments both ways on whether to impute raw data or to impute transformed data. You may choose either.

Hints:

- Don’t assume non-missing data is actually not missing.
- Use `print = FALSE` in your mice call to suppress that annoying output.

Hints for extra credit:

- Using `gather` may again be handy for getting the data in a format `ggplot` likes (then again, it might not, depending on your solution).
- To ensure certain points are “on top” of the other points, you can add a second `geom_whatever` which operates on a subsetted data. If this occurs later in the `ggplot` stack, it will plot on top.
- The `scales = "free"` argument in `facet_` may be useful (or `scales = "free_x"` for just “x”, same for “y”).

Exercise 3

Use the `testthat` package to add unit tests to the `getMovieData` function from the first assignment. You may use your function from that assignment, or my solution if you prefer.

Important If you run into any failing tests (indicating a bug), you *are not required to fix the code*, however, do identify what the bug or oversight might be.

Notes:

- For whatever reason, placing `expect_*` directly in RMarkdown doesn’t work, even passing chunk option `error = TRUE`. To get around this, wrap all your `expect_*` statements in a `test_that` call as well as passing `error = TRUE` as a chunk option. I will explain when we talk about testing inside packages the actual purpose of the `test_that` function.

```
```${r, error = TRUE}``  
test_that(desc = "my tests", {
 expect_*(...)
 expect_*(...)
})
````
```

- I don’t want to put any sort of limitation like “You must have at least k tests”, but I would be surprised if anyone managed to provide complete code coverage with less than 20-30 tests. For reference, I wrote 45 tests, although I tend to be a bit overzealous in testing. On the other hand, I did it rapidly so I’m sure I didn’t test every corner.
- In case I forget to mention it in class, as opposed to the normal programming principle of DRY (Don’t Repeat Yourself) where you should put any code you’ll use more than once into a function, in testing, you Do Repeat Yourself.

Exercise 4

Ensure your system is prepared for R package development! Install the necessary tools:

- Windows: RTools
- Mac: XCode (You may be able to just run `xcode-select --install` at the command line instead)
- Linux: install `r-base-dev` package from `apt-get`, etc.

Test this making sure this returns `TRUE`:

```
library(devtools)  
has_devel()
```

```
## '/Library/Frameworks/R.framework/Resources/bin/R' --no-site-file \  
## --no-environ --no-save --no-restore --quiet CMD SHLIB foo.c  
##
```

```
## [1] TRUE
```