Lab 12 - Math 58B: Building Multiple Linear Regression Models

due Tuesday April 28, 2026

your name here

```
library(tidyverse)
library(broom)
library(praise)
```

Lab Goals

The goal of the lab is to put into practice the tools we've been using to analyze multiple linear models. In particular, we will:

- a full analysis including multiple variables
- including inference in the linear model
- checking technical conditions using residual plots
- choosing variables for the multiple linear regression

Getting started

For formatting the the linear model output, use the **broom** package. Also, **tidyverse** will continue to be used for data wrangling and `ggplot()` for data visualization.

For using data wrangling verbs, you might see: https://r4ds.had.co.nz/transform.html (remembering the verbs `filter()`, `arrange()`, `select()`, `mutate()`, `summarize()`, and `group_by()`.)

For `ggplot()`, look for inspiration on this cheat sheet:

https://raw.githubusercontent.com/rstudio/cheatsheets/main/data-visualization.pdf

Load packages & data

The data is on student evaluations of 94 professors from the University of Texas at Austin, from the **openintro** package.

```
library(tidyverse)
library(openintro)
data(evals)
```

Grading the professor

Many college courses conclude by giving students the opportunity to evaluate the course and the instructor anonymously. However, the use of these student evaluations as an indicator of course quality and teaching effectiveness is often criticized because these measures may reflect the influence of non-teaching related characteristics, such as the physical appearance of the instructor. The research found that instructors who are viewed to be better looking receive higher instructional ratings.[1]

The data

The data were gathered from end of semester student evaluations for a large sample of professors at the University of Texas, Austin. Additionally, six students rated the professors' physical appearance.[2] In the resulting data frame, each row represents a different course and columns represent variables about the courses and professors.

variable	description
`score`	average professor evaluation score: (1) very unsatisfactory - (5) excellent.
`rank`	rank of professor: teaching, tenure track, tenured.
`ethnicity`	ethnicity of professor: not minority, minority.
`gender`	gender of professor: female, male.
`language`	language of school where professor received education: English or non-English.
`age`	age of professor.
`cls_perc_eval`	percent of students in class who completed evaluation.
`cls_did_eval`	number of students in class who completed evaluation.

[1] Daniel S. Hamermesh, Amy Parker, Beauty in the classroom: instructors pulchritude and putative pedagogical productivity, *Economics of Education Review*, Volume 24, Issue 4, August 2005, Pages 369-376, ISSN 0272-7757, 10.1016/j.econedurev.2004.07.013. http://www.sciencedirect.com/science/article/pii/S0272775704001165.

[2] Data are from (slightly modified) *Data Analysis Using Regression and Multilevel/Hierarchical Models* (Gelman & Hill, 2007).

variable	description
`cls_students`	total number of students in class.
`cls_level`	class level: lower, upper.
`cls_profs`	number of professors teaching sections in course in sample: single, multiple.
`cls_credits`	number of credits of class: one credit (lab, PE, etc.), multi-credit.
`bty_f1lower`	beauty rating of professor from lower level female: (1) lowest - (10) highest.
`bty_f1upper`	beauty rating of professor from upper level female: (1) lowest - (10) highest.
`bty_f2upper`	beauty rating of professor from second upper level female: (1) lowest - (10) highest.
`bty_m1lower`	beauty rating of professor from lower level male: (1) lowest - (10) highest.
`bty_m1upper`	beauty rating of professor from upper level male: (1) lowest - (10) highest.
`bty_m2upper`	beauty rating of professor from second upper level male: (1) lowest - (10) highest.
`bty_avg`	average beauty rating of professor.
`pic_outfit`	outfit of professor in picture: not formal, formal.
`pic_color`	color of professor's picture: color, black & white.

Exploring the data

Q.

What are the observational units in this study?

Q.

Is this an observational study or an experiment? The original research question posed in the paper is whether beauty leads directly to the differences in course evaluations. Given the study design, is it possible to answer this question as it is phrased? If not, rephrase the question.

Q.

Describe the distribution of `score`. Is the distribution skewed? What does that tell you about how students rate courses? Is this what you expected to see? Why, or why not?

Q.

Excluding `score`, select two other variables and describe their relationship using an appropriate visualization (scatterplot, side-by-side boxplots, barplot, or histogram).

Simple linear regression

The fundamental phenomenon suggested by the study is that better looking teachers are evaluated more favorably. Let's create a scatterplot to see if this appears to be the case:

```
ggplot(evals) +
  geom_point(aes(x = bty_avg, y = score))
```

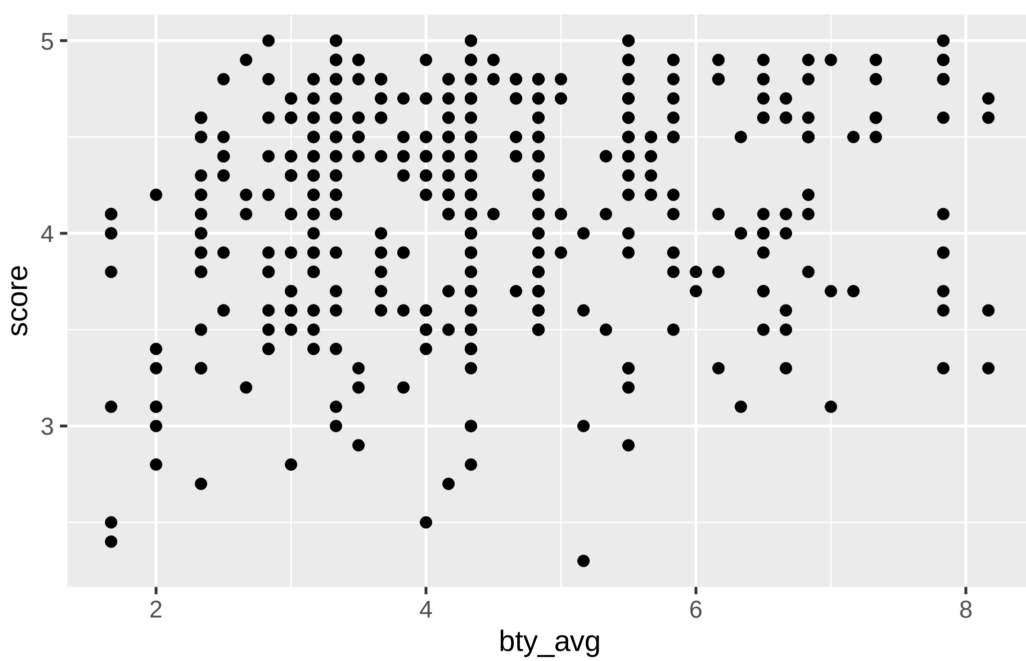

Before we draw conclusions about the trend, compare the number of observations in the data frame with the approximate number of points on the scatterplot. Is anything awry?

Q.

Replot the scatterplot, but this time use the layer `geom_jitter()`. What was misleading about the initial scatterplot?

Q.

Let's see if the apparent trend in the plot is something more than natural variation. Fit a linear model called `m_bty` to predict average professor score by average beauty rating and add the line to your plot using `geom_smooth(method = "lm")`. Write out the equation for the linear model and interpret the slope. Is average beauty score a statistically discernible predictor? Does it appear to be a meaningful predictor?

Q.

Use residual plots to evaluate whether the conditions of least squares regression are reasonable. Provide plots and comments on whether the linear model seems reasonable here.

Note that the `augment()` function will give you residuals (`.resid`) as well as predicted values (`.fitted`). Once you have the observations, you can plot them using `ggplot()`.

Multiple linear regression

What about the model that includes `bty_avg` and `gender`? Does the residual plot meet the technical conditions?

```
m_bty_gen <- lm(score ~ bty_avg + gender, data = evals)
m_bty_gen |> tidy()
```

```
# A tibble: 3 x 5
  term        estimate std.error statistic   p.value
  <chr>          <dbl>     <dbl>     <dbl>     <dbl>
1 (Intercept)   3.75      0.0847     44.3  6.23e-168
2 bty_avg       0.0742    0.0163      4.56 6.48e-  6
3 gendermale    0.172     0.0502      3.43 6.52e-  4
```

Q.

P-values and parameter estimates should only be trusted if the conditions for the regression are reasonable. Verify that the conditions for this model are reasonable using a residual plot. Make a residual plot to assess the technical conditions (remember to pipe the linear model into the `augment()` function).

Q.

With `gender` in the model, is `bty_avg` still a discernible predictor of `score`? Has the addition of `gender` to the model changed the parameter estimate for `bty_avg`? (Find the model and pipe it into the `tidy()` output.)

Q.

What is the equation of the line corresponding to males? (*Hint:* For males, the parameter estimate is multiplied by 1.) For two professors who received the same beauty rating, which gender tends to have the higher course evaluation score?

Q.

Create a new model predicting score by multiplying `bty_avg` and `gender`. Write down (separately) the the equation of the line corresponding to females and males? (*Hint:* For males, the parameter estimate is multiplied by 1.) You will have two equations. Do your two equations have different slopes? Do the equations of the line match the plot above? Multiplying two variables together produces what we call *interaction*.

```
m_bty_gen_int <- lm(score ~ bty_avg * gender, data = evals)
m_bty_gen_int |> tidy()
```

```
# A tibble: 4 x 5
  term                estimate std.error statistic   p.value
  <chr>                  <dbl>     <dbl>     <dbl>     <dbl>
1 (Intercept)           3.95      0.118      33.5  2.92e-125
2 bty_avg               0.0306    0.0240      1.28 2.02e-  1
3 gendermale           -0.184     0.153      -1.20 2.32e-  1
4 bty_avg:gendermale    0.0796    0.0325      2.45 1.46e-  2
```

Q.

Create a new model called `m_bty_rank` with `gender` removed and `rank` added in. How does R appear to handle categorical variables that have more than two levels? Note that the rank variable has three levels: `teaching`, `tenure track`, `tenured`. Write down the three separate models for teach rank. Make a scatterplot to suggest whether `bty_avg` seem to interact. If so, run an interaction model.

To Turn In

The search for the best model

So far, we've considered `bty_avg`, `gender`, and `rank` as variables which might predict `score`. For the next part of the lab, we'll consider the following questions that will help us build a single linear regression model on multiple explanatory / predictor variables. How can we go about finding which variables to use in a model? Which are the best variables? Which are the most important variables? Which are the variables that are not discernible?

Start with a full model that predicts professor score based on rank, ethnicity, gender, language of the university where they got their degree, age, proportion of students that filled out evaluations, class size, course level, number of professors, number of credits, average beauty rating, outfit, and picture color.

Q1. Which variable is worst?

Which variable would you expect to have the highest p-value in this model? Why? *Hint:* Think about which variable would you expect to not have any association with the professor score. (Use your instincts, not anything technical about the data.)

Check your suspicions by running the model. Include the model output in your response.

Q2. Ethnicity

Interpret the coefficient associated with the ethnicity variable.

Q3. Pare down the model

a. Drop the variable with the highest p-value and re-fit the model. Did the coefficients and discernibility of the other explanatory variables change? (One of the things that makes multiple regression interesting is that coefficient estimates depend on the other variables that are included in the model.) If not, what does this say about whether or not the dropped variable was collinear with the other explanatory variables?
b. Should we keep going and remove other variables? Which one(s)? Remove any variables (one at a time) that do not seem to be important to predicting `score`.

Q4. Technical Conditions

Verify that the conditions for this model are reasonable using a residual plots.

Q5. Model Interpretation

Based on your final model, describe the characteristics of a professor and course at University of Texas at Austin that would be associated with a high evaluation score.

Q6. Independence technical condition?

The original paper describes how these data were gathered by taking a sample of professors from the University of Texas at Austin and including all courses that they have taught. Considering that each row represents a course, could this new information have an impact on any of the conditions of linear regression?

Q7. Infer to what population?

Would you be comfortable generalizing your conclusions to apply to professors generally (at any university)? Why or why not?

```r
praise()
```

```
[1] "You are posh!"
```