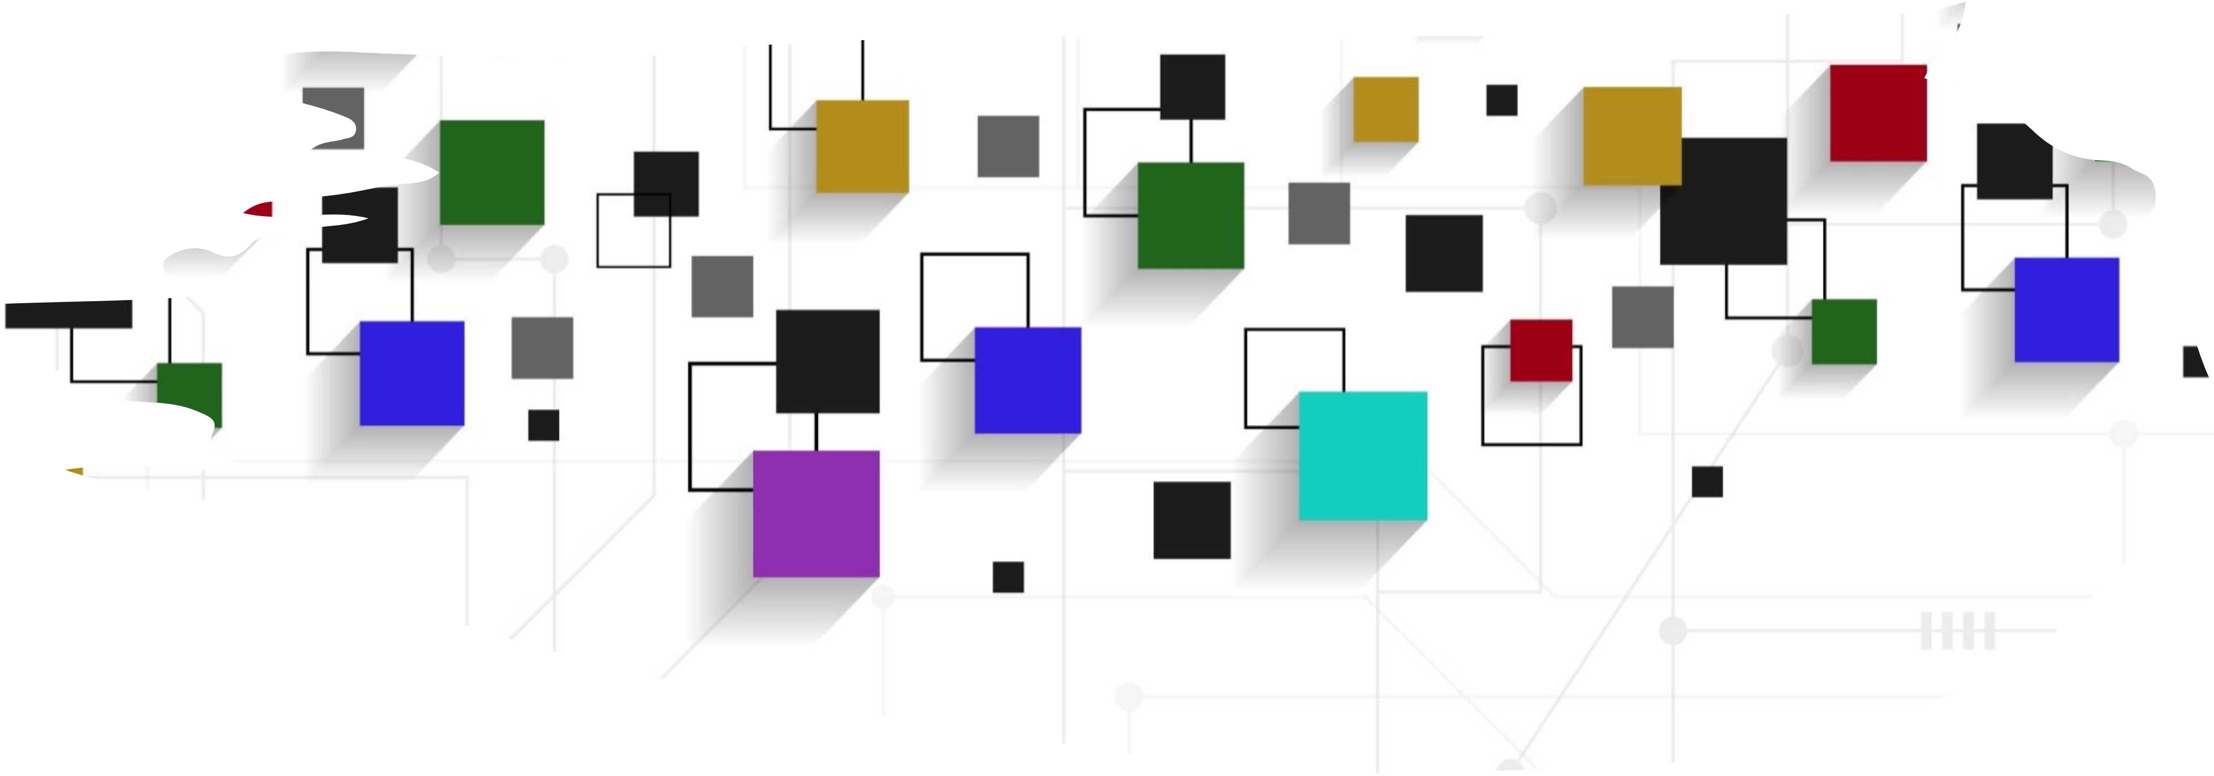

CogLab: Data Collection / Inferences

WEEK 12

sona data collection plan

- pre-register on aspredicted.org
- add disqualifiers to your study
- add survey code to cognition.run study URL in Sona
- modify cognition.run initJsPsych
- send approval request to Donna
- post 100 timeslots, Nov 30 deadline
- start working on project analyses!

11	Thursday, November 9, 2023	Weeks 11-13: Data Collection
12	Tuesday, November 14, 2023	Data Collection continued…
12	Thursday, November 16, 2023	**Psychonomics Conference: NO CLASS**
12	Sunday, November 19, 2023	**Formative Assignment (R Inferential) Due**
13	Tuesday, November 21, 2023	Data Collection continued…
13	Thursday, November 23, 2023	**THANKSGIVING BREAK!!! NO CLASS**
14	Tuesday, November 28, 2023	W14: Odds and Ends
14	Wednesday, November 29, 2023	**Project Milestone #7 (Analyses) Due**
14	Thursday, November 30, 2023	W14 continued…
14	Sunday, December 3, 2023	**Project Milestone #8 (Poster Draft) Due**
15	Tuesday, December 5, 2023	W15: Wrapping Up
15	Thursday, December 7, 2023	**Project Milestone #9 (Poster Symposium) Due**
16	Sunday, December 17, 2023	**Project Milestone #10 (Final Report) Due**

disqualifiers

Disqualifiers

Participants must not have completed or have a pending sign-up for ANY of these studies:

My Studies | All Studies

search...

- Assistance game (online) (Inactive)
- Block game
- Connector Word Game!
- Semantic Association Task
- Semantic Integration Task
- Sentence Content Judgement (Inactive)
- Sentence Experiment (Inactive)
- Sentence Judgements (Inactive)
- Sentence Object Recognition (Inactive)

Available

Selected

changing study URL

Study URL

https://xz4vnidz3d.cognition.run?sona_id=%SURVEY_CODE%

If the text `%SURVEY_CODE%` is included in the URL, the system will replace that with a unique code for the participant, to make it easier to identify who completed the study. You can also configure it so that participants receive credit in the system immediately after finishing the survey. If you are using Cognition, add `?sona_id=%SURVEY_CODE%` to the end of the URL to make use of this feature.

Detailed Help

copy cognition finish URL

Website

View Study Website

Sample Link with Embedded ID Code

Cognition Finish URL

"https://bowdoin.sona-systems.com/webstudy_credit.aspx?ex

Instructions

You can also configure it so that participants receive credit in the system immediately after finishing the survey. If you are using Cognition, add `?sona_id=%SURVEY_CODE%` to the end of the URL to make use of this feature.

Detailed Help

change initJsPsych within cognition.run

```
const jsPsych = initJsPsych({
    show_progress_bar: true,
    auto_update_progress_bar: false,

    on_finish: function(data){

      let sona_id = jsPsych.data.urlVariables()['sona_id']
      window.location.assign("https://bowdoin.sona-systems.com/webstudy_credit.aspx?experiment_id=218&credit_token=ecd1a56cc7194604ba296e8e927f57ac&survey_code=" +sona_id)
    }

    });
```

send for approval

- attach [approved IRB protocol]

Study Status	Not visible to participants : Not Approved Send Request Active study : Does not appear on list of available studies -- must also be approved Online (web) study : Administered outside the system

add timeslot

Study Menu

- View/Administer Time Slots
- Timeslot Usage Summary
- Download Participant List
- Contact Participants
- View Bulk Mail Summary
- Change Study Information
- Participant Study View
- Study Modification Log
- Copy Study
- Delete Study

Add Timeslots : Semantic Association Task

This study was created as an online (web) study. Because a participant may participate in an online study at any time, most researchers create a single timeslot. The single timeslot contains the maximum number of participants who may participate, and has a final participation date of the last date that participants may participate.

NOTE: You are adding timeslots to a study that is unapproved, so participants will not be able to sign up for the study.

Final Participation Date	Sunday, November 12, 2023
Final Participation Time	9:00 AM
Max. Number of Participants	1

Add This Timeslot

recap: Nov 7/9, 2023

- what we covered:
  - linear regression, t-tests, and ANOVAs
- your to-do's were:
  - *resubmit:* formative assignment #2
  - *finalize:* experiment
  - submit: pre-registration

today’s agenda

- linear regression continued
- two-way/multiple linear regression

linear model: assumptions

- "all models are wrong, but some are useful" (Box, 1976)
- the model does not know where the data come from or whether they are appropriate for the model; that is your responsibility as a researcher
  - linearity
  - normality of residuals
  - homoskedasticity
  - independence of observations

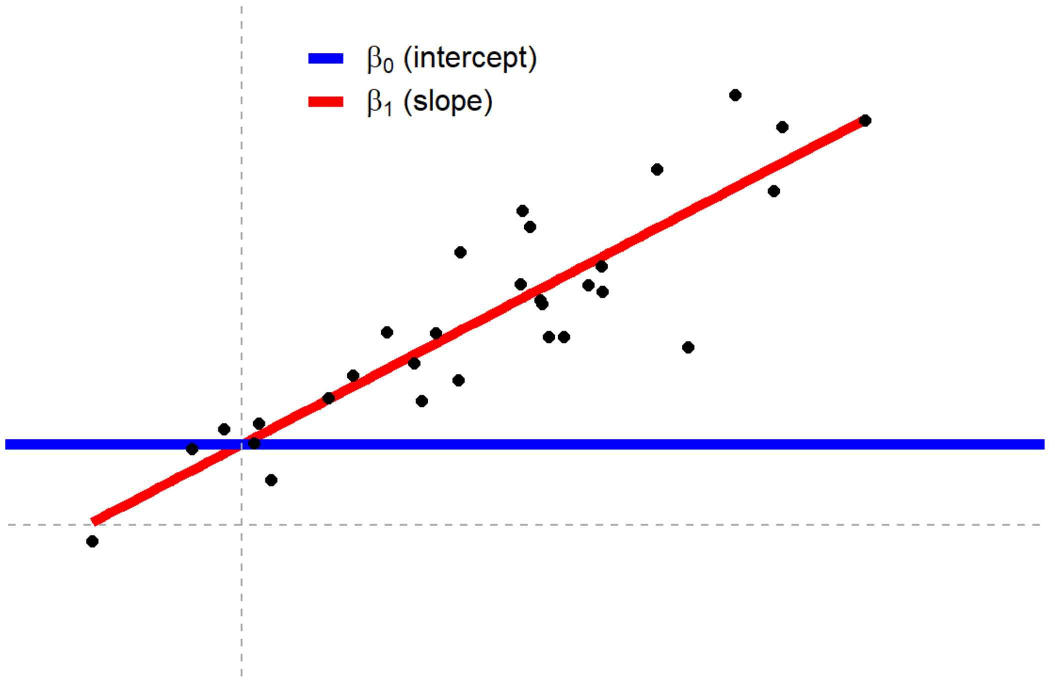

https://lindeloev.github.io/tests-as-linear/

inspecting the model

```r
install.packages("performance", dependencies = TRUE)
install.packages("see", dependencies = TRUE)
install.packages("patchwork", dependencies = TRUE)
```

```r
library(performance)
check_model(full_iris_model)
```

- first we install the performance, see, and patchwork packages
- load performance
- check the model
- minor variations are ok, major variations are warnings!

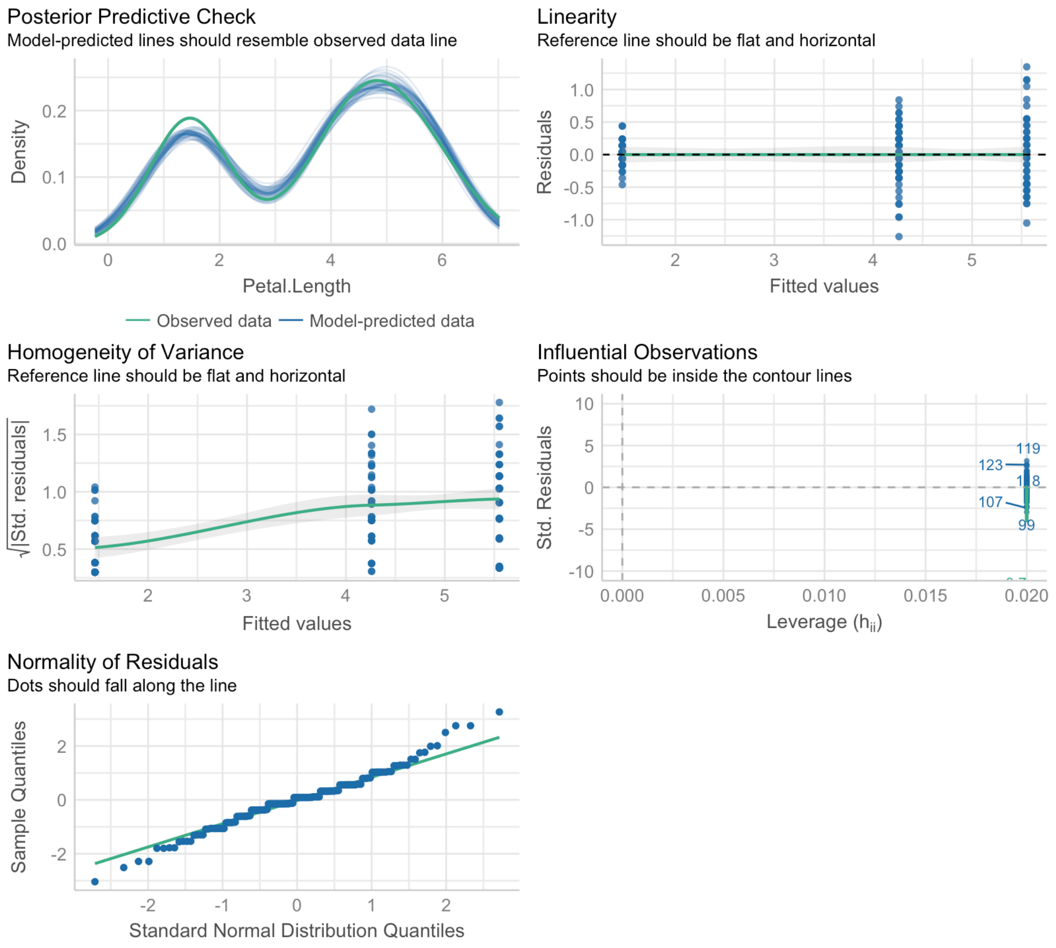

multiple linear regression

- often, we want to look at the influence of more than one variable on our response measures
- a multiple linear regression is a model that attempts to find the relationship between a dependent variable and more than one independent variable
  - $Y = aX_1 + bX_2 + c$
  - Y: dependent variable
  - $X_{1,2}$: independent variables

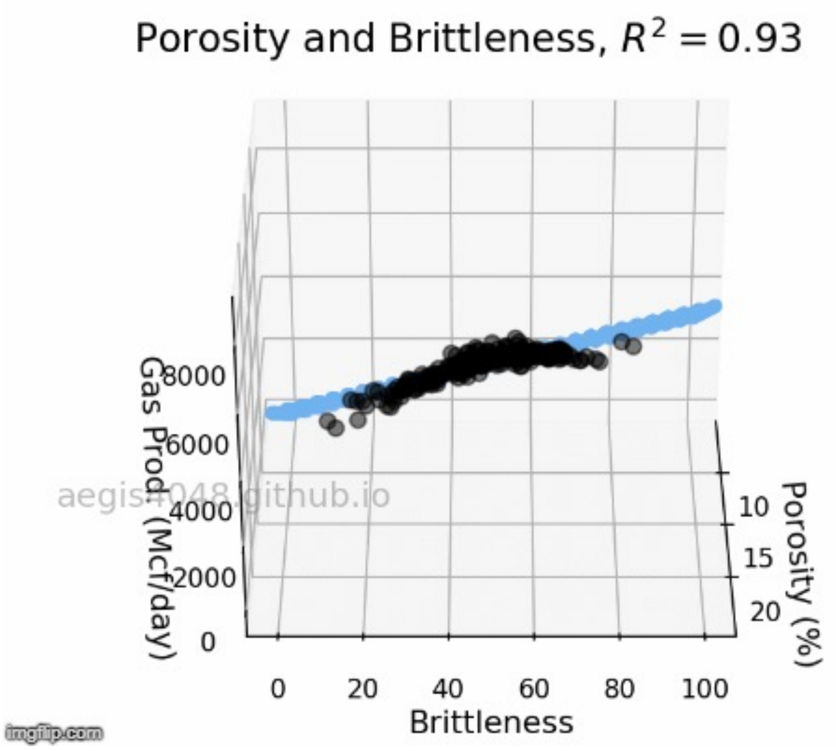

multiple linear regression: data

- we will use the jobsatisfaction dataset from the datarium package
- install the package datarium
- new heading (# multiple linear regression) & code chunk
- load and view the jobsatisfaction dataset

```
data("jobsatisfaction", package = "datarium")
View(jobsatisfaction)
```

id	gender	education_level	score
1	male	school	5.51
2	male	school	5.65
3	male	school	5.07
4	male	school	5.51
5	male	school	5.94
6	male	school	5.80
7	male	school	5.22
8	male	school	5.36
9	male	school	4.78
10	male	college	6.01
11	male	college	6.01
12	male	college	6.45

multiple linear regression: exploration

- let's explore the data:
    - visualize the pattern via a boxplot

multiple linear regression: exploration

- let's explore the data:
  - visualize the pattern via a boxplot
  - do you see differences in job satisfaction?

```
jobsatisfaction %>%
  ggplot()+
  geom_boxplot(aes(x = gender, y = score, color = education_level))
```

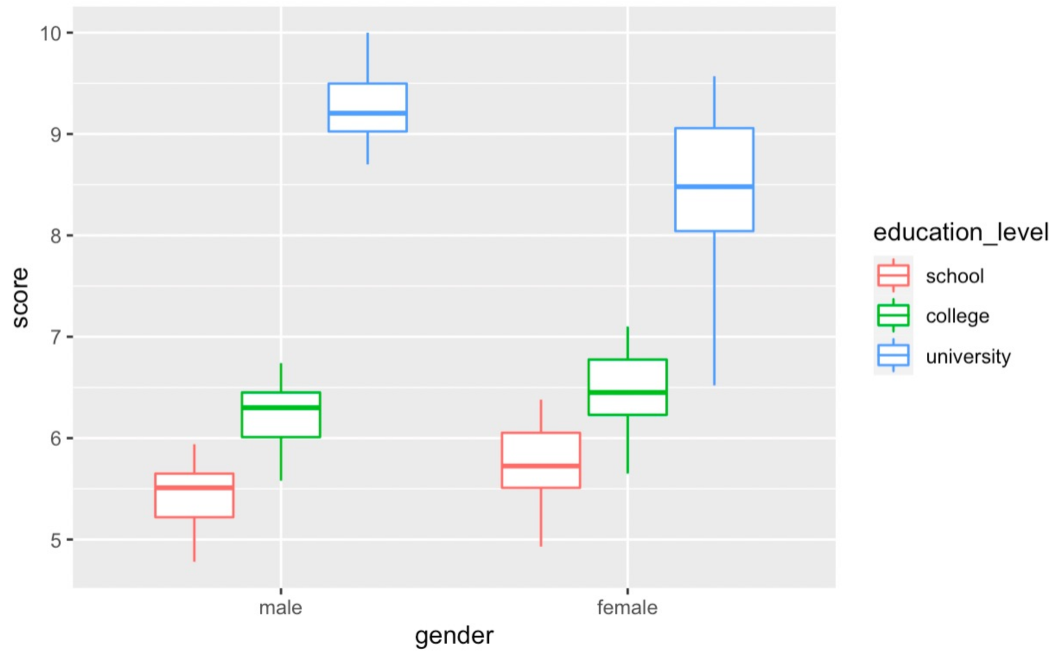

multiple linear regression: research question

- does job satisfaction vary as a function of gender and education level?
- dependent variable?
- independent variable?

```
jobsatisfaction %>%
  ggplot()+
  geom_boxplot(aes(x = gender, y = score, color = education_level))
```

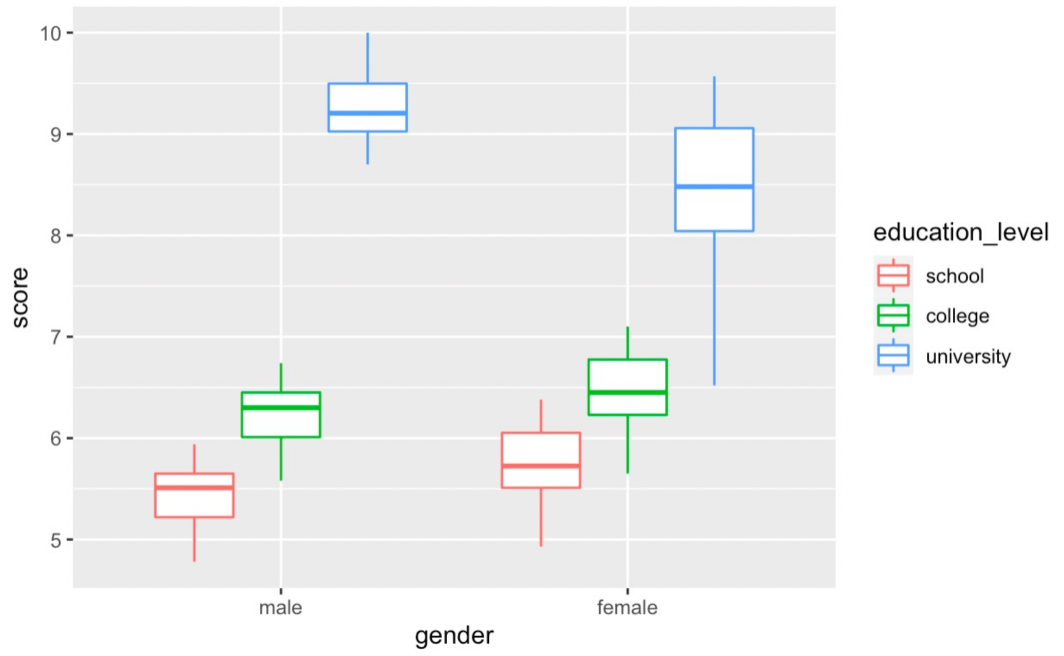

main effects

- when you have multiple variables in your experiment design, there are few different possibilities for how the pattern of data might look
- you could have the dependent variable vary as a function of IV1 and/or IV2 (main effects), and these effects might interact with each other
- main effects refer to differences in means of levels of an independent variable
- what is an example of a main effect for the jobsatisfaction dataset?
- what would the plot of this main effect look like?

id	gender	education_level	score
1	male	school	5.51
2	male	school	5.65
3	male	school	5.07
4	male	school	5.51
5	male	school	5.94
6	male	school	5.80
7	male	school	5.22
8	male	school	5.36
9	male	school	4.78
10	male	college	6.01
11	male	college	6.01
12	male	college	6.45

interactions

- interactions refer to situations when the difference in means between IV1's levels differs based on the levels of IV2, i.e., you cannot simply infer a difference in means
- what is an example of an interaction for the jobsatisfaction dataset?
- what would the plot of this interaction look like?

id	gender	education_level	score
1	male	school	5.51
2	male	school	5.65
3	male	school	5.07
4	male	school	5.51
5	male	school	5.94
6	male	school	5.80
7	male	school	5.22
8	male	school	5.36
9	male	school	4.78
10	male	college	6.01
11	male	college	6.01
12	male	college	6.45

visually…

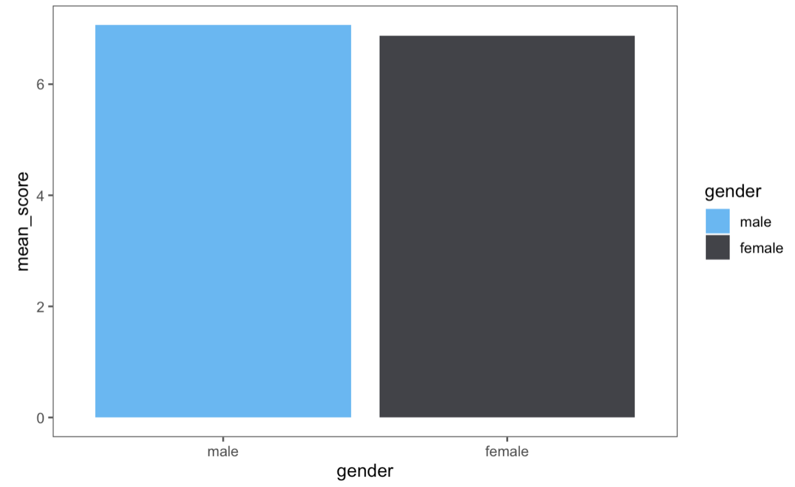
mean_score
0
2
4
6
male
female
gender
gender
male
female

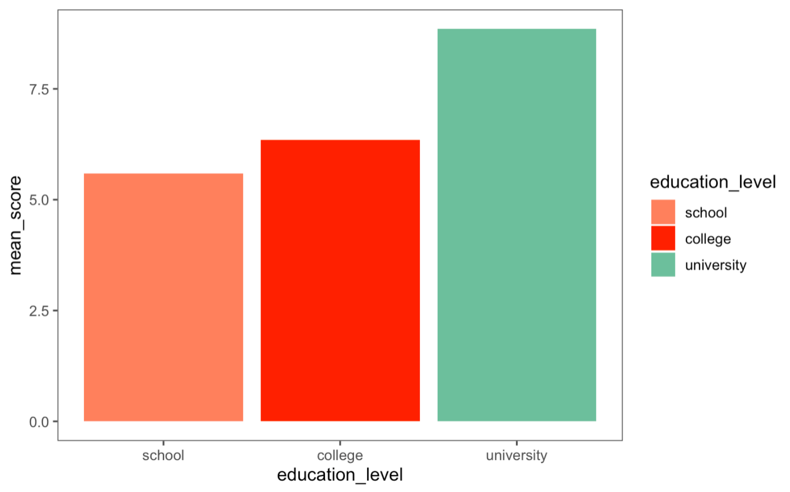
mean_score
0.0
2.5
5.0
7.5
school
college
university
education_level
education_level
school
college
university

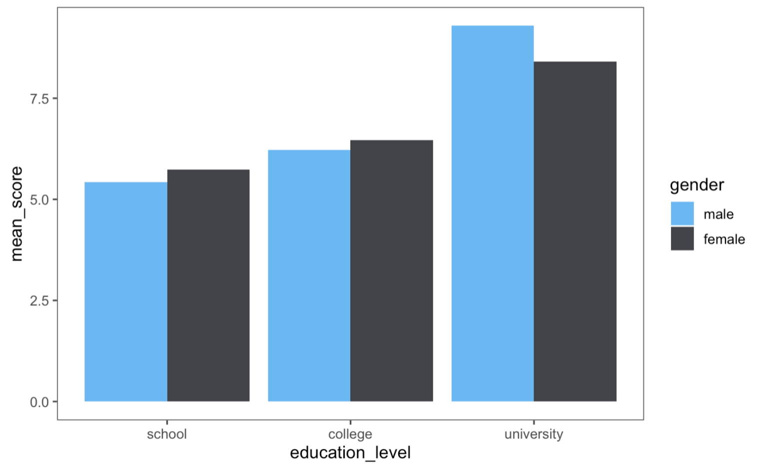
mean_score
0.0
2.5
5.0
7.5
school
college
university
education_level
gender
male
female

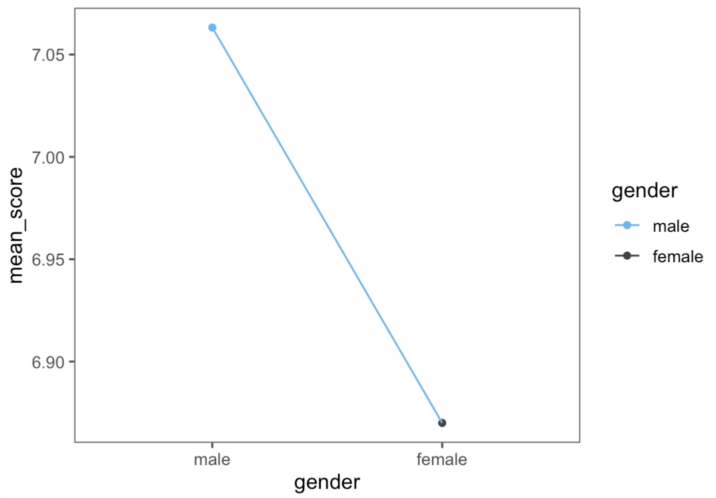
mean_score
6.90
6.95
7.00
7.05
male
female
gender
gender
male
female

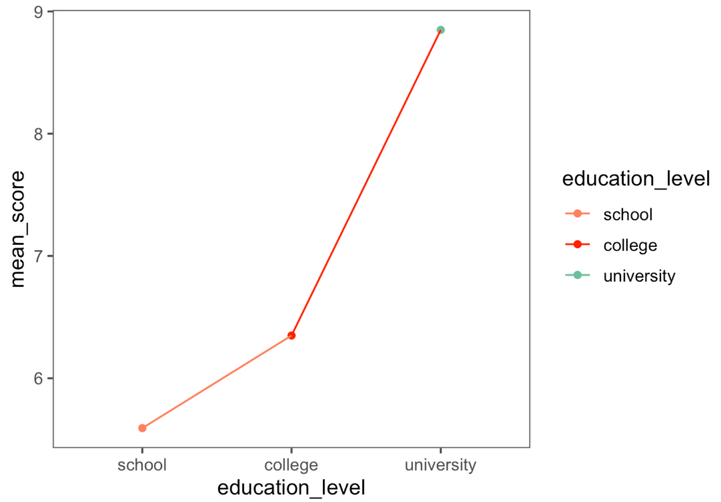
mean_score
6
7
8
9
school
college
university
education_level
education_level
school
college
university

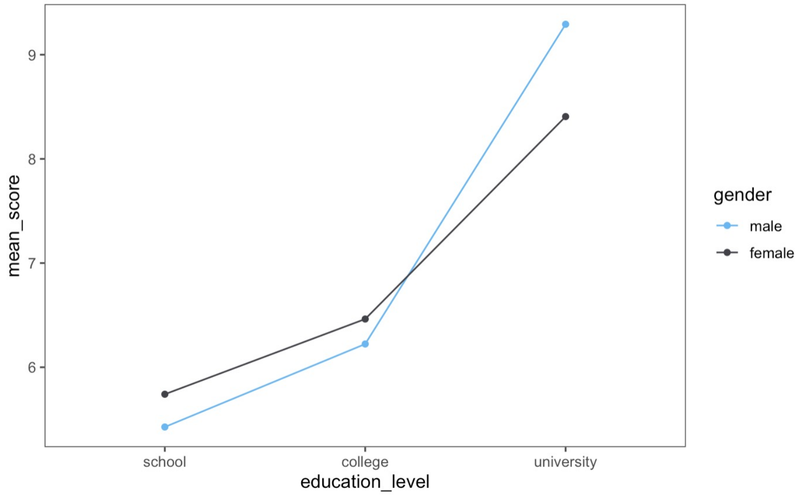
mean_score
6
7
8
9
school
college
university
education_level
gender
male
female

mathematically…

- main effect of gender:
  - mean (male) – mean (female)
- main effect of education level
  - mean(school) – mean (college)
  - mean(college) – mean (university)
  - mean(university) – mean(school)
- interaction (difference of differences)
  - diff(male-female)$_{school}$– diff(male-female)$_{college}$
  - diff(male-female)$_{university}$– diff(male-female)$_{college}$
  - diff(male-female)$_{school}$– diff(male-female)$_{university}$

gender <fctr>	education_level <fctr>	mean <dbl>	sd <dbl>
male	school	5.426667	0.3638681
male	college	6.223333	0.3396322
male	university	9.292000	0.4445422
female	school	5.741000	0.4744225
female	college	6.463000	0.4746941
female	university	8.406000	0.9379078

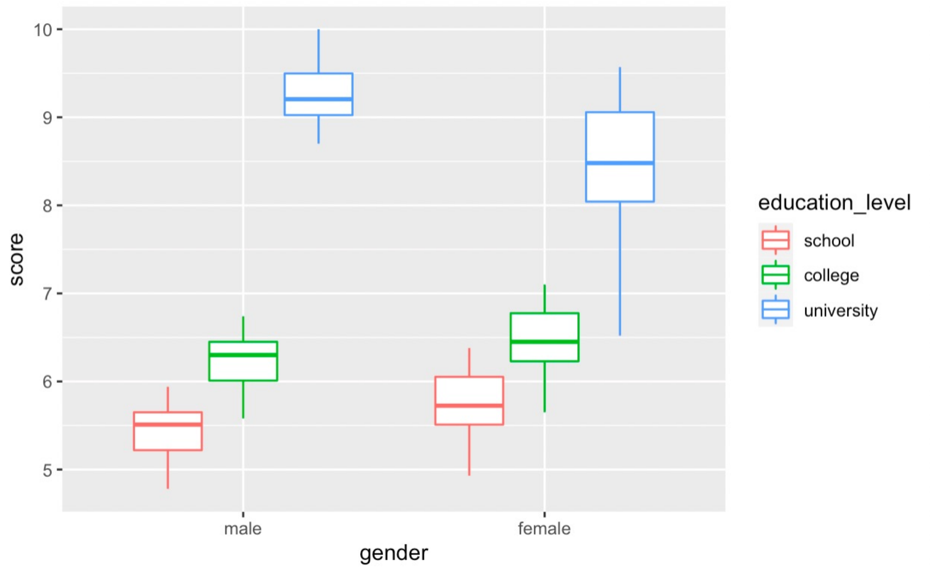

multiple linear regression in R

- we define a job_model that uses a linear model as before, with separate terms for main effects and interactions
- how do we view the results of this model?

```
job_model = lm(data = jobsatisfaction,
               score ~ gender + education_level + gender:education_level)
```

```
summary(job_model)
car::Anova(job_model)
```

interpreting multiple regression outputs

- the intercept is always the reference condition, and all estimates are relative to this reference condition
- in this case, the reference category is male, school

```
Call:
lm(formula = score ~ gender + education_level + gender:education_level,
    data = jobsatisfaction)

Residuals:
    Min      1Q  Median      3Q     Max
-1.8860 -0.2325 -0.0040  0.3297  1.1640

Coefficients:
                                      Estimate Std. Error t value Pr(>|t|)
(Intercept)                            5.42667    0.18334  29.599  < 2e-16 ***
genderfemale                           0.31433    0.25272   1.244  0.21915
education_levelcollege                 0.79667    0.25928   3.073  0.00337 **
education_leveluniversity              3.86533    0.25272  15.295  < 2e-16 ***
genderfemale:education_levelcollege   -0.07467    0.35740  -0.209  0.83533
genderfemale:education_leveluniversity -1.20033   0.35266  -3.404  0.00129 **
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 0.55 on 52 degrees of freedom
Multiple R-squared:  0.8829,    Adjusted R-squared:  0.8717
F-statistic: 78.45 on 5 and 52 DF,  p-value: < 2.2e-16
```

linear regression and ANOVAs

- you just conducted an ANOVA!
- ANOVAs are special cases of linear regression models, when the predictors are *categorical*
- two-way ANOVA equation
  - $y=\beta_0+ \beta_1X_1+ \beta_2X_2 + \beta_3X_1X_2$
  - note that the X's here are different independent variables
  - H_0: $\beta_1=0$ (for X_1 main effect)
  - H_0: $\beta_2=0$ (for X_2 main effect)
  - H_0: $\beta_3=0$ (for interaction)

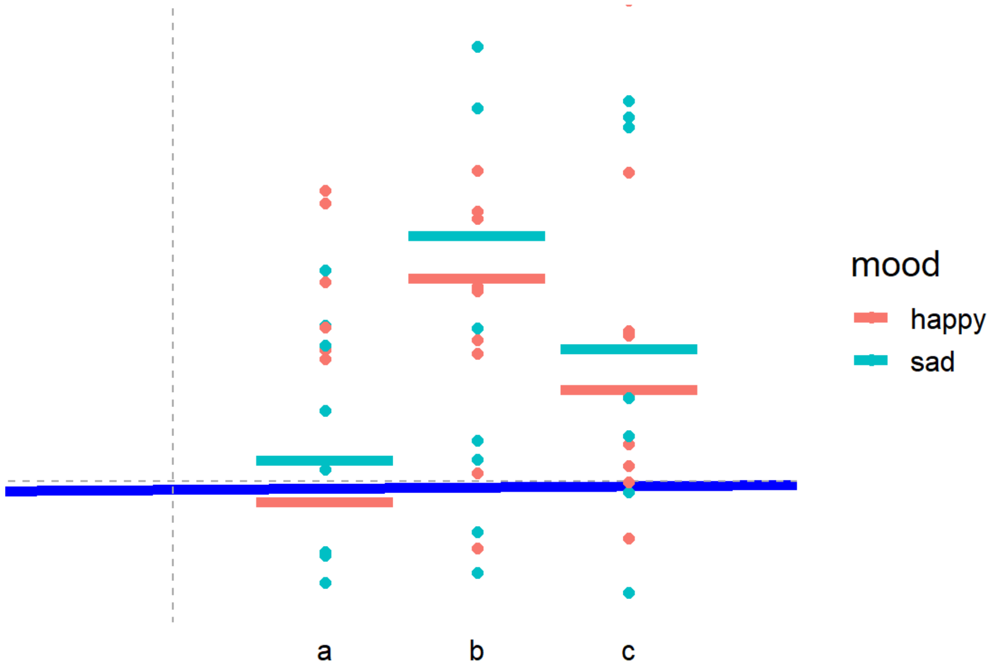

https://lindeloev.github.io/tests-as-linear/

understanding main effects & interactions

- in the case of categorical IVs, viewing the car::Anova() result is useful to understand broad patterns
- we see a main effect of education level, but it is qualified by the interaction with gender
- what does this mean?

```
> car::Anova(job_model)
Anova Table (Type II tests)

Response: score
                          Sum Sq Df  F value    Pr(>F)
gender                     0.225  1   0.7447  0.392115
education_level          113.684  2 187.8921 < 2.2e-16 ***
gender:education_level     4.440  2   7.3379  0.001559 **
Residuals                 15.731 52
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1
```

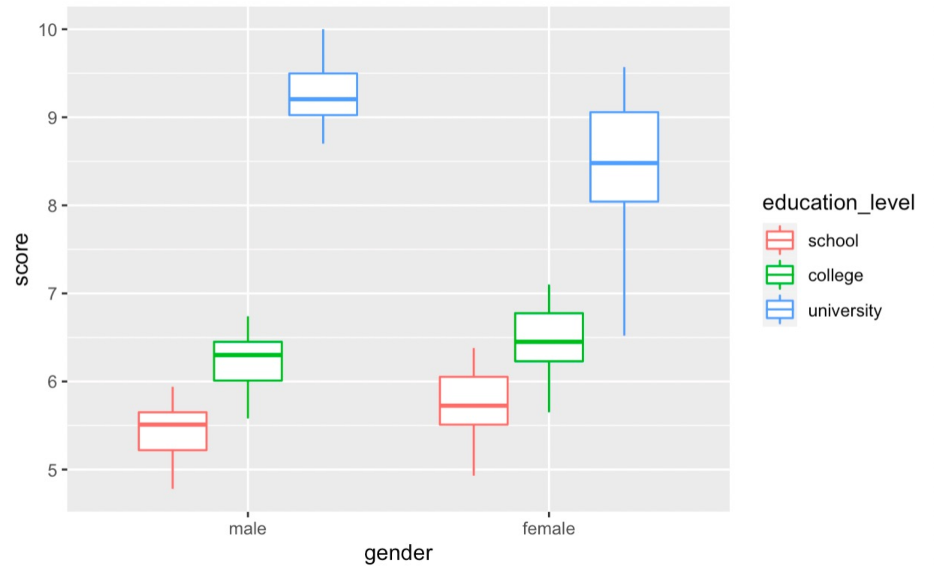

decomposing an interaction

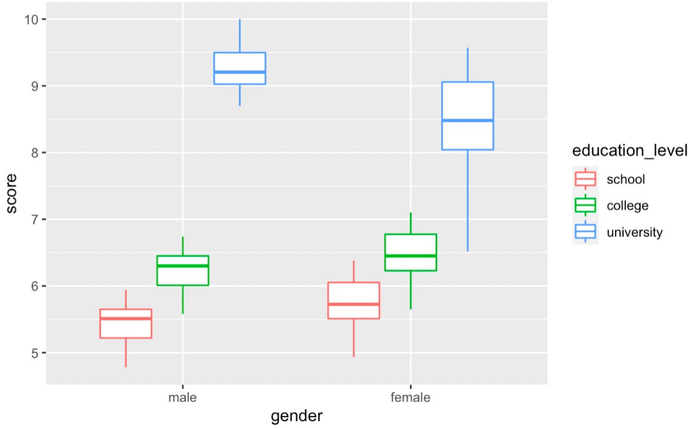

- when there is a significant interaction, we want to go in and understand the nature of this interaction
- interaction (difference of differences)
  - diff(male-female)$_{school}$– diff(male-female)$_{college}$
  - diff(male-female)$_{university}$– diff(male-female)$_{college}$
  - diff(male-female)$_{school}$– diff(male-female)$_{university}$
- where do we think the difference may be?

using emmeans

- we use emmeans as before, except now we specify a conditional effect
- what do the contrasts tell us?

```r
emmeans::emmeans(job_model,
                 pairwise ~ gender | education_level,
                 adjust="tukey")
```

```
$emmeans
education_level = school:
 gender emmean    SE df lower.CL upper.CL
 male     5.43 0.183 52     5.06     5.79
 female   5.74 0.174 52     5.39     6.09

education_level = college:
 gender emmean    SE df lower.CL upper.CL
 male     6.22 0.183 52     5.86     6.59
 female   6.46 0.174 52     6.11     6.81

education_level = university:
 gender emmean    SE df lower.CL upper.CL
 male     9.29 0.174 52     8.94     9.64
 female   8.41 0.174 52     8.06     8.76

Confidence level used: 0.95

$contrasts
education_level = school:
 contrast      estimate    SE df t.ratio p.value
 male - female   -0.314 0.253 52  -1.244  0.2191

education_level = college:
 contrast      estimate    SE df t.ratio p.value
 male - female   -0.240 0.253 52  -0.948  0.3473

education_level = university:
 contrast      estimate    SE df t.ratio p.value
 male - female    0.886 0.246 52   3.602  0.0007
```

assumptions: multiple linear regression

- same as before
- linearity
- normality of residuals
- homoskedasticity
- independence of observations

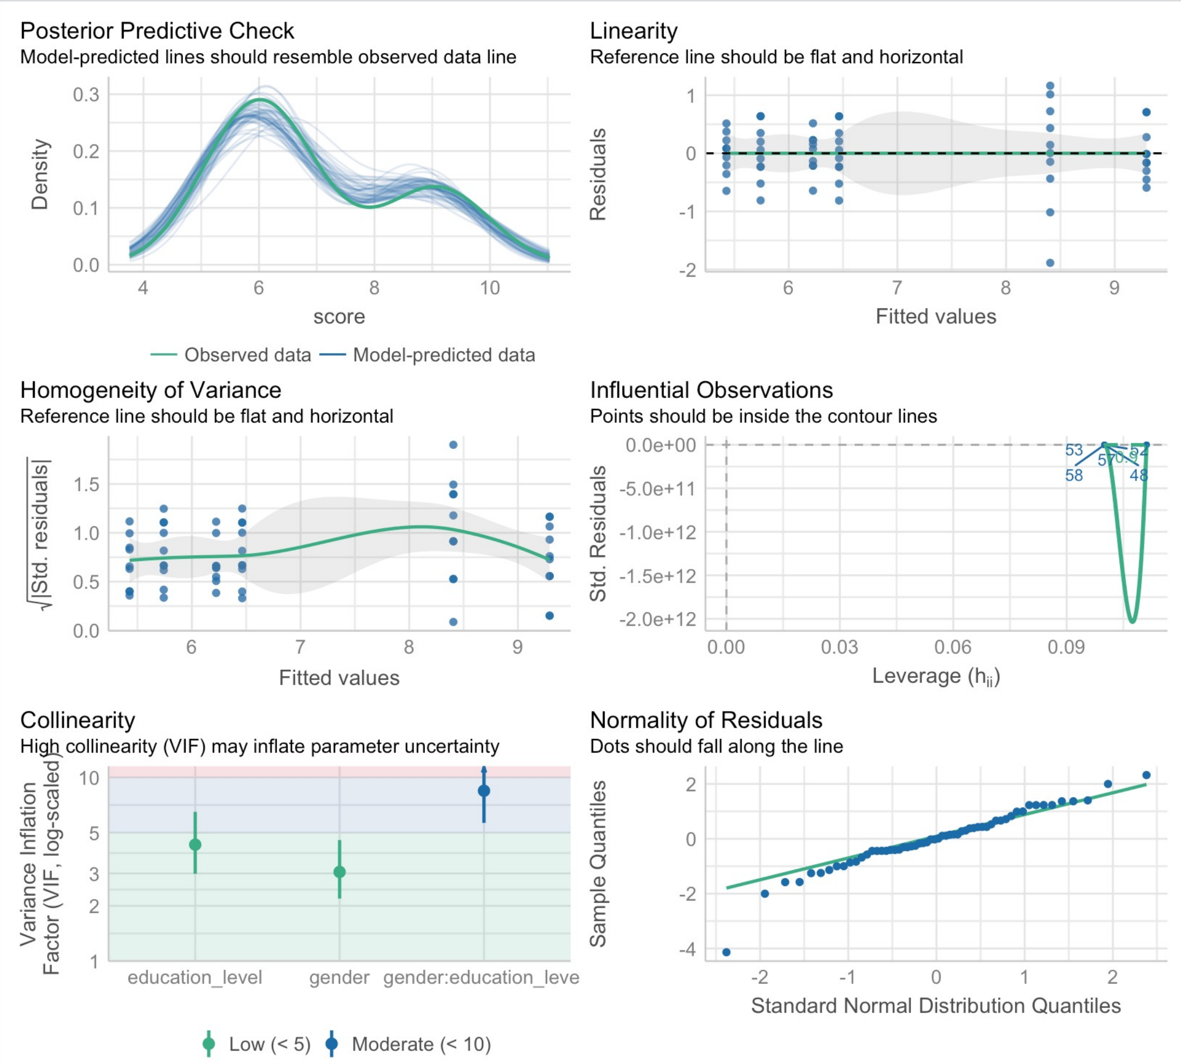

revisiting class data

- run all chunks
- view the priming data
- what are the IVs?
- what is the DV?
- double check data types for IV/DV

first-R-notebook.Rmd × | priming_data × | jobsatisfaction ×

Filter

	ID	rt	relatedness	prime	response	type	correct	block_number	target	correct_key
1	896801386	221	related	geck	a	shared	TRUE	1	horse	A
2	896801386	292	unrelated	geck	a	shared	TRUE	1	apple	A
3	399427091	647	unrelated	foobly	l	direct	TRUE	1	horse	L
4	399427091	640	unrelated	foobly	a	direct	TRUE	1	horse	A
5	399427091	562	related	dodish	l	direct	TRUE	1	horse	L
6	399427091	487	related	foobly	l	direct	TRUE	1	apple	L
7	399427091	663	related	mipp	a	shared	TRUE	1	apple	A
8	399427091	674	unrelated	mipp	a	shared	TRUE	1	horse	A

```
priming_data       1473 obs. of 10 variables
  $ ID          : num [1:1473] 8.97e+08 8.97e+08 3.99e+08 3.99e+08 3.99e+08 ...
  $ rt          : num [1:1473] 221 292 647 640 562 487 663 674 747 582 ...
  $ relatedness : Factor w/ 3 levels "novel","related",..: 2 3 3 3 2 2 2 3 3 3 .
  $ prime       : chr [1:1473] "geck" "geck" "foobly" "foobly" ...
  $ response    : chr [1:1473] "a" "a" "l" "a" ...
  $ type        : Factor w/ 3 levels "direct","novel",..: 3 3 1 1 1 1 3 3 3 1 ..
  $ correct     : chr [1:1473] "TRUE" "TRUE" "TRUE" "TRUE" ...
  $ block_number: chr [1:1473] "1" "1" "1" "1" ...
  $ target      : chr [1:1473] "horse" "apple" "horse" "horse" ...
  $ correct_key : chr [1:1473] "A" "A" "L" "A" ...
```

multiple linear regression

- run a multiple linear regression
- examine the assumptions plot
  - linearity
  - normality of residuals
  - homoskedasticity
  - independence of observations

```
rt_lm_model = lm(data = priming_data, rt ~ relatedness + type + relatedness:type)
summary(rt_lm_model)
```

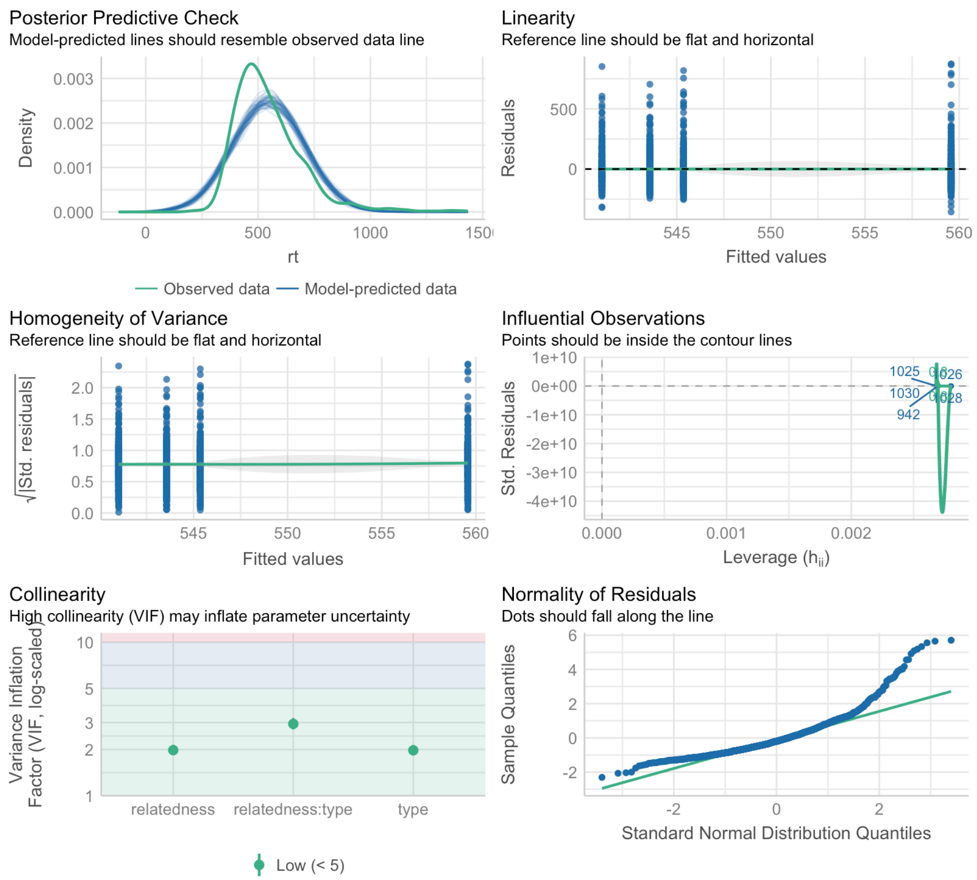

non-independent designs

- whenever multiple observations are collected from participants, especially in within-subject designs, we cannot use a typical linear model / ANOVA
- usual solution: repeated measures ANOVA on the means per condition per ID
- problem: we lose information when we aggregate data before analysis

ANOVA: limitations

- limited to continuous DVs
  - examples of non-continuous/categorical DVs?
  - changes the distribution of your variables, violates the normality assumption
  - common distributions: binomal (yes/no, correct/incorrect), multinomial (know,don't know, TOT, other), poisson (counting number of website visitors)
- limited to categorical IVs
  - examples of continuous IVs?
- cannot deal with missing data
- cannot handle nested/clustered design
  - male/females in sectors in cities
  - trials in subjects in conditions
- cannot handle unbalanced design
  - different number of trials (after exclusion) for each subject?

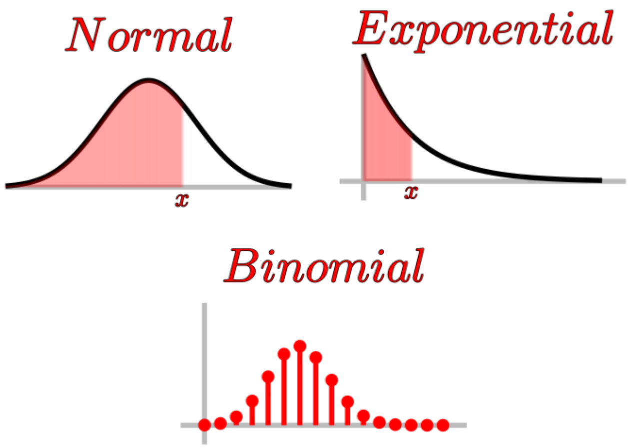

a flexible model

- linear/generalized mixed effects models!
- these models consider the variability due to:
  - missing data
  - categorical/continuous IVs and DVs
  - unbalanced designs
  - clustered designs (no collapsing into means)
- think of them as the parent models from which special cases such as t-tests and ANOVAs are derived
- different 'lines/curves' are fit for each individual and for each item, with their own slope and intercept, instead of "averaging" across everyone

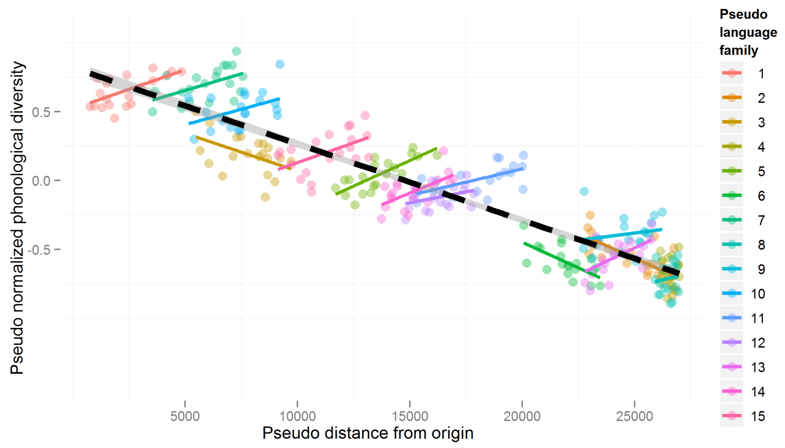

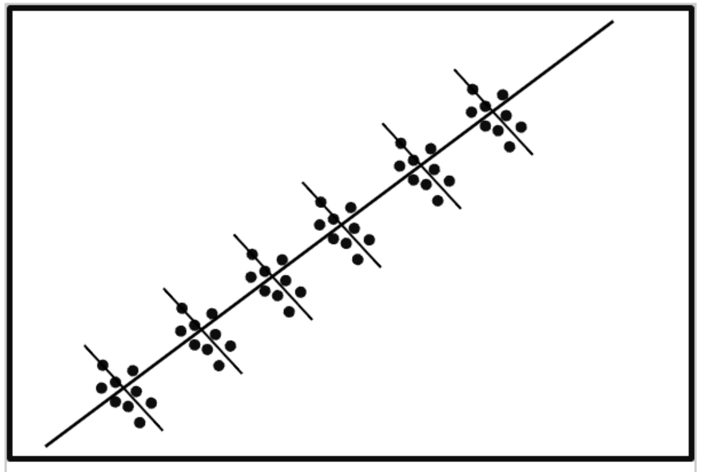

mixed linear model in R

- similar format, just specifying the within-subject variable as the “random” part of the model

```r
library(lmerTest)
rt_model = lmer(data = priming_data,
                rt ~ relatedness*type + (1|ID))

summary(rt_model)
```

```
Linear mixed model fit by REML. t-tests use Satterthwaite's method ['lmerModLmerTest']
Formula: rt ~ relatedness * type + (1 | ID)
   Data: priming_data

REML criterion at convergence: 18587.2

Scaled residuals: 
    Min      1Q  Median      3Q     Max 
-2.2600 -0.5893 -0.1396  0.3975  5.8610 

Random effects:
 Groups   Name        Variance Std.Dev.
 ID       (Intercept)  8049     89.72  
 Residual             17024    130.48  
Number of obs: 1473, groups:  ID, 27

Fixed effects:
                                Estimate Std. Error       df t value Pr(>|t|)    
(Intercept)                      536.528     18.721   29.152  28.659   <2e-16 ***
relatednessunrelated              16.707      9.580 1441.289   1.744   0.0814 .  
typeshared                        -1.730      9.566 1441.696  -0.181   0.8566    
relatednessunrelated:typeshared  -15.259     13.613 1441.522  -1.121   0.2625    
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

Correlation of Fixed Effects:
            (Intr) rltdns typshr
rltdnssnrlt -0.256              
typeshared  -0.258  0.500       
rltdnssnrl:  0.181 -0.704 -0.702
```

main effects and interactions

- car::Anova()

```
car::Anova(rt_model)
```

```
> car::Anova(rt_model)
Analysis of Deviance Table (Type II Wald chisquare tests)

Response: rt
                  Chisq Df Pr(>Chisq)
relatedness      1.8072  1     0.1788
type             1.8490  1     0.1739
relatedness:type 1.2565  1     0.2623
```

assumptions check

- same as before

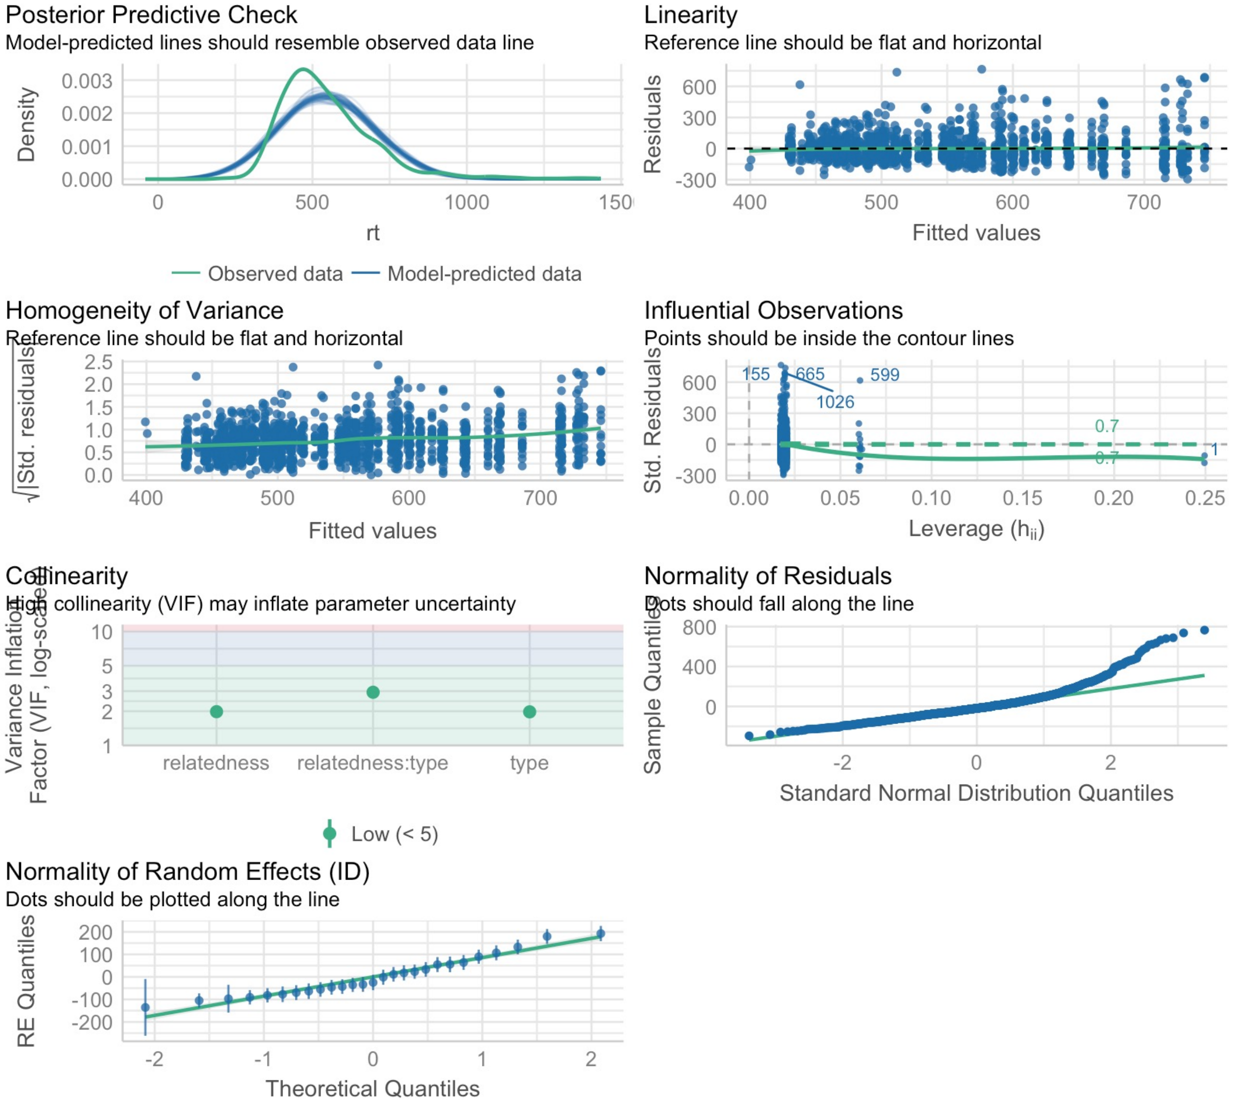

big takeaways

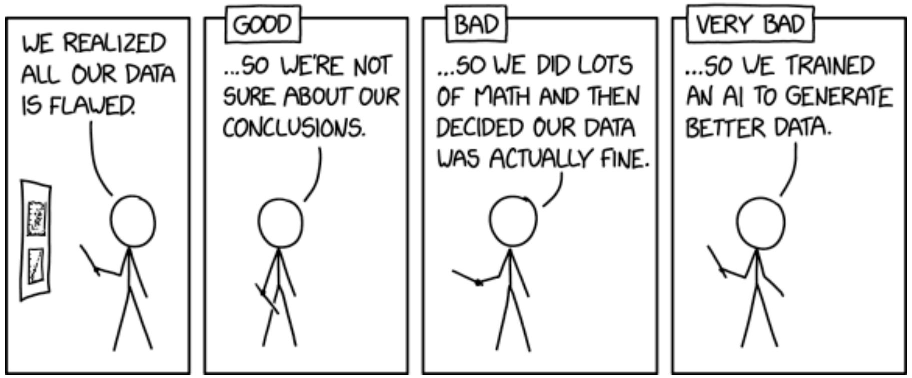

- statistical analyses are often taught from the framework of different-tests-for-different-data
- but…the same principles underlie most tests you encounter
- there now exist more robust, flexible methods that overcome the limitations of t-tests and ANOVAs and allow you to truly capture different levels of variability
- there ALSO exist methods of analysis that do not heavily rely on p-values (frequentist statistics) and account for prior information in making inferences (Bayesian statistics)
- keep an open mind and try to find connections between methods you read about and see around you!

next time

- **before** class
  - *monitor:* data collection on Sona
  - *submit:* formative assignment #3 (due Nov 19)
  - *work on:* project milestone #7 (analyses, due Nov 29)
- **during** class (Nov 21)
  - analysis review
  - prolific data collection
  - poster design