PSYC402-week-17-LME-1

Rob Davies (Lancaster University)

PSYC402 – Rob Davies – week 17

Figure 1: flickr: Stefan 'Take off into 2017'

Targets for Week 17 – Ideas and skills, and foundations for later work

1. concepts: multilevel data and multilevel modeling
2. skills: visualization – examine overall and within-class trends
3. skills: run linear models over all data – and within each class
4. skills: use the lmer() function to fit models of multilevel data

This week, we concentrate on understanding *multilevel data*, and some of the ideas behind the practice of multilevel modeling

- *What, why, when*

This week, we concentrate on understanding *multilevel data*, and some of the ideas behind the practice of multilevel modeling

- *What, why, when*
  - *What, why* – What are multilevel models? – Why do we use them?

This week, we concentrate on understanding *multilevel data*, and some of the ideas behind the practice of multilevel modeling

- *What, why, when*
  - *What, why* – What are multilevel models? – Why do we use them?
  - *When* should they be applied?

This week, we concentrate on understanding *multilevel data*, and some of the ideas behind the practice of multilevel modeling

- *What, why, when*
  - *What, why* – What are multilevel models? – Why do we use them?
  - *When* should they be applied?
- *How* do we conduct multilevel models? – we get started this week, develop skills over next weeks

This week, we concentrate on understanding *multilevel data*, and some of the ideas behind the practice of multilevel modeling

- *What, why, when*
  - *What, why* – What are multilevel models? – Why do we use them?
  - *When* should they be applied?
- *How* do we conduct multilevel models? – we get started this week, develop skills over next weeks
- Next week, we will start talking about Linear Mixed-effects models

Phenomena and data sets in the social sciences often have a multilevel structure

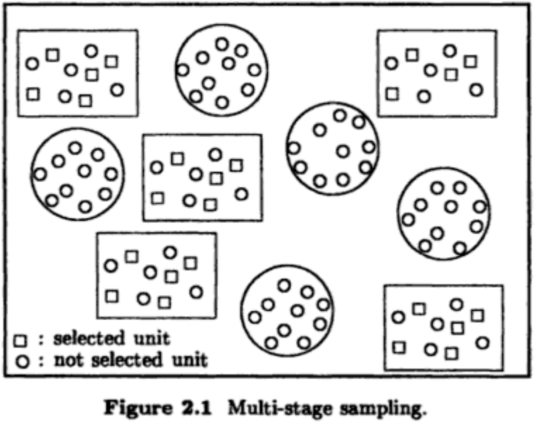

Figure 2: Snijders and Bosker (2012) Multistage sample

Repeated measures or clustered data samples

- Test the same people multiple times

Repeated measures or clustered data samples

- Test the same people multiple times
  - Pre- and post-treatment

Repeated measures or clustered data samples

- Test the same people multiple times
  - Pre- and post-treatment
  - Multiple stimuli – everyone sees the same stimuli

Repeated measures or clustered data samples

- Test the same people multiple times
  - Pre- and post-treatment
  - Multiple stimuli – everyone sees the same stimuli
  - Repeated testing – follow learning, development within individuals – in longitudinal designs

Repeated measures or clustered data samples

- Test the same people multiple times
  - Pre- and post-treatment
  - Multiple stimuli – everyone sees the same stimuli
  - Repeated testing – follow learning, development within individuals – in longitudinal designs
- Do multi-stage sampling

Repeated measures or clustered data samples

- Test the same people multiple times
  - Pre- and post-treatment
  - Multiple stimuli – everyone sees the same stimuli
  - Repeated testing – follow learning, development within individuals – in longitudinal designs
- Do multi-stage sampling
  - Find (sample) classes or schools – test (sample) children within classes or schools

Repeated measures or clustered data samples

- Test the same people multiple times
  - Pre- and post-treatment
  - Multiple stimuli – everyone sees the same stimuli
  - Repeated testing – follow learning, development within individuals – in longitudinal designs
- Do multi-stage sampling
  - Find (sample) classes or schools – test (sample) children within classes or schools
  - Find (sample) clinics – test (sample) patients within clinics

The key insight: observations are clustered – correlated – not independent

Hierarchical – multilevel – structures are everywhere but they have typically been ignored – at a price

- Traditional analytic methods have typically required the researcher to aggregate their data
  - Example: averaging the responses made by a participant to different stimuli

Hierarchical – multilevel – structures are everywhere but they have typically been ignored – at a price

- Traditional analytic methods have typically required the researcher to aggregate their data
  - Example: averaging the responses made by a participant to different stimuli
- Or to ignore the hierarchical structure in their data
  - Example: analysing the responses made by some pupils while ignoring the fact that the pupils were tested in different classes

What is the price?

- Ignoring structure

What is the price?

- Ignoring structure
  - can mean analyses are less sensitive because they do not fully account for random differences

What is the price?

- Ignoring structure
  - can mean analyses are less sensitive because they do not fully account for random differences
  - can mean increased chance of detecting effects spuriously because detected effects due to error variation

Next, we look at an example that makes our concerns concrete

When – first example: children nested within classes

- If you test participants belonging to different groups, e.g., observe the grades of children recruited from different classes in a school

When – first example: children nested within classes

- If you test participants belonging to different groups, e.g., observe the grades of children recruited from different classes in a school
  - The scores for the participants are observations that occupy the lowest level of a hierarchy

When – first example: children nested within classes

- If you test participants belonging to different groups, e.g., observe the grades of children recruited from different classes in a school
  - The scores for the participants are observations that occupy the lowest level of a hierarchy
  - Those observations are understood to be nested within higher level sampling units – the children's scores are nested within classes

First example – pupils' English language scores in classes in a school in Brazil

- We analyze the end-of-year school subject grades for a sample of 292 children studied by Golino and Gomes (2014) in Brazil

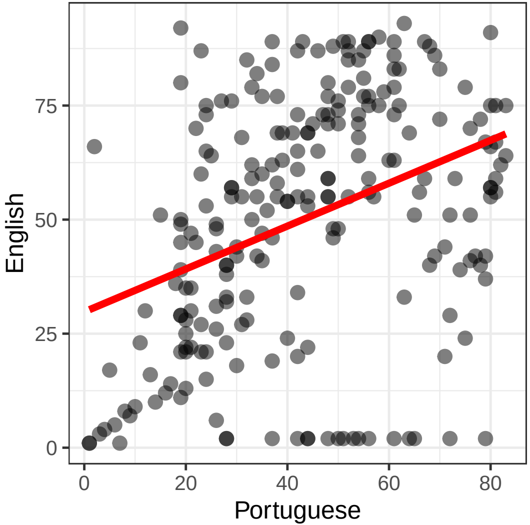

First example – pupils' English language scores in classes in a school in Brazil

- We analyze the end-of-year school subject grades for a sample of 292 children studied by Golino and Gomes (2014) in Brazil
- A scatterplot shows that if a child has a higher grade in Portuguese she will tend to have a higher grade in English

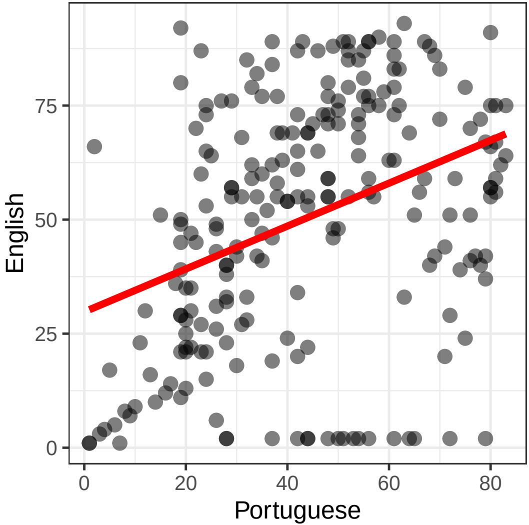

We can estimate a linear model that takes children's English grades as the outcome (dependent) variable and their Portuguese grades as the predictor (independent) variable

```
lm(english ~ portuguese, data = BAFACALO_DATASET)
```

The R code expresses a linear model

$$y = \beta_0 + \beta_1 X + e \quad (1)$$

Our linear model predicts a child's English grade

$$y = \beta_0 + \beta_1 X + e \tag{2}$$

- y – English grade for each child

Our linear model predicts a child's English grade

$$y = \beta_0 + \beta_1 X + e \tag{2}$$

- y – English grade for each child
- β_0 – intercept

Our linear model predicts a child's English grade

$$y = \beta_0 + \beta_1 X + e \tag{2}$$

- y – English grade for each child
- β_0 – intercept
- $\beta_1 X$ – estimated effect β_1, the difference in English grade, associated with differences in X the Portuguese grade, for each child

Our linear model predicts a child's English grade

$$y = \beta_0 + \beta_1 X + e \tag{2}$$

- y – English grade for each child
- β_0 – intercept
- $\beta_1 X$ – estimated effect β_1, the difference in English grade, associated with differences in X the Portuguese grade, for each child
- e – differences between the observed English grade and the English grade predicted by the relationship with Portuguese grades, for each child

Linear model yields the estimate that for unit increase in Portuguese grade there is an associated .5 increase in English grade, on average

```
##
## Call:
## lm(formula = english ~ portuguese, data = BAFACALO_DATASET)
##
## Residuals:
##     Min      1Q  Median      3Q     Max
## -64.909 -17.573   2.782  20.042  53.292
##
## Coefficients:
##             Estimate Std. Error t value Pr(>|t|)
## (Intercept) 29.77780    3.81426   7.807 2.11e-13 ***
## portuguese   0.47001    0.07897   5.952 9.91e-09 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 25.02 on 228 degrees of freedom
##   (62 observations deleted due to missingness)
## Multiple R-squared:  0.1345, Adjusted R-squared:  0.1307
## F-statistic: 35.43 on 1 and 228 DF,  p-value: 9.906e-09
```

The linear model ignores the higher-level structure – the distinction between classes: does that matter?

- A scatterplot shows that the relationship between grades in Portuguese and grades in English *does* vary between classes

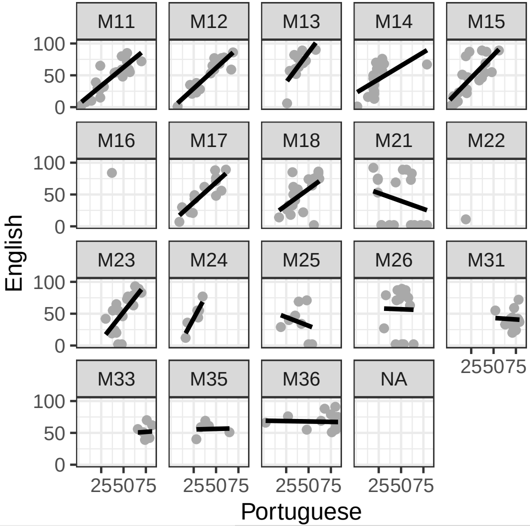

The linear model ignores the higher-level structure – the distinction between classes: does that matter?

- A scatterplot shows that the relationship between grades in Portuguese and grades in English *does* vary between classes
- Here, we split the plot to show the relationship between Portuguese and English grades for each child – separately for each class

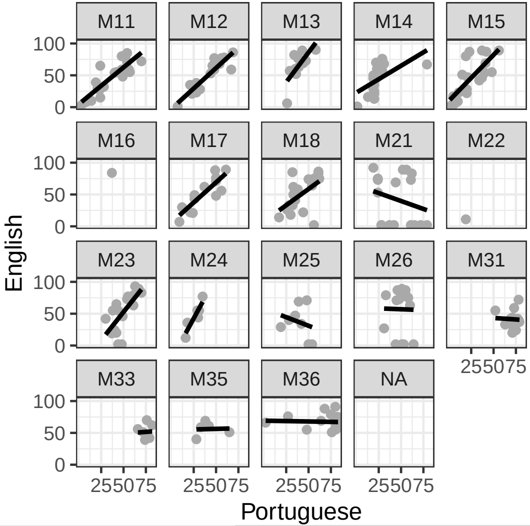

We can deal with the differences between groups by fitting a separate model for each class

- We can analyze the relationship between Portuguese and English grades separately for each class

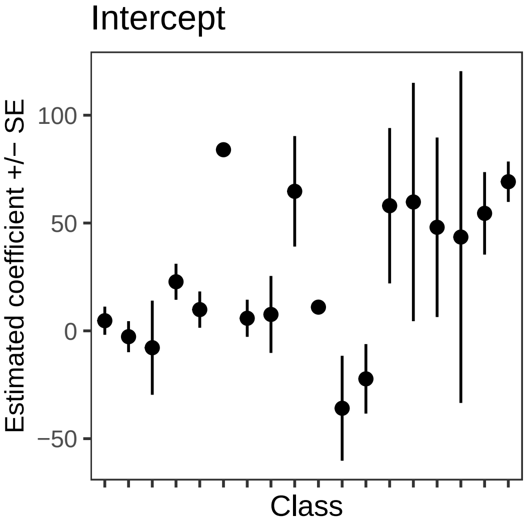

We can deal with the differences between groups by fitting a separate model for each class

- We can analyze the relationship between Portuguese and English grades separately for each class
- We can then extract the estimated coefficients β_0, β_1 and plot them

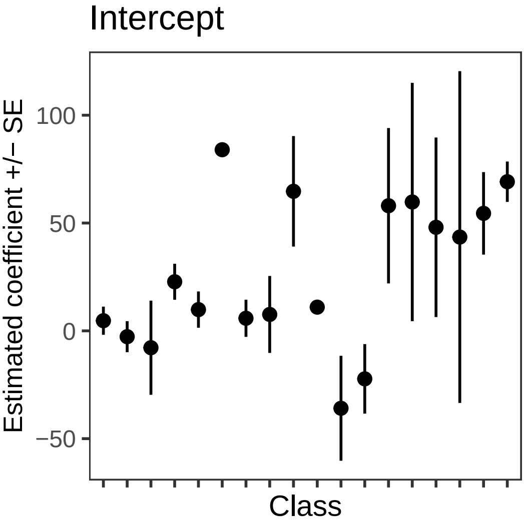

We can deal with the differences between groups by fitting a separate model for each class

- We can analyze the relationship between Portuguese and English grades separately for each class
- We can then extract the estimated coefficients β_0, β_1 and plot them
- Here we see how intercepts vary between classes

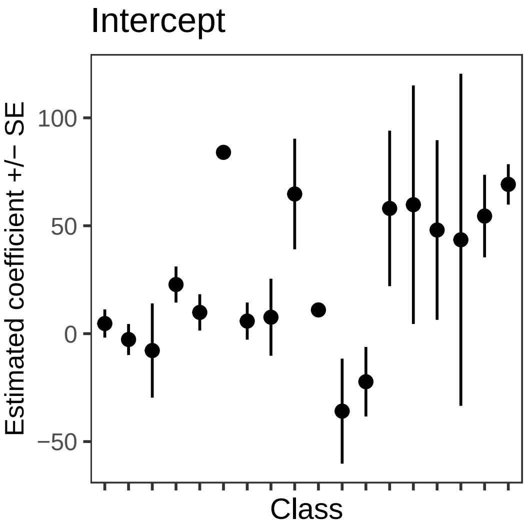

We can deal with the differences between groups by fitting a separate model for each class

- Here we see how slopes of the "effect" of Portuguese on English varies between classes

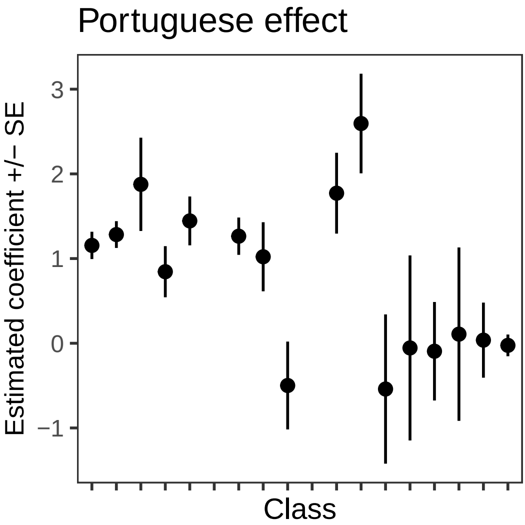

Next we look at one method for analyzing multilevel data which has important limitations: then we look at a better method – **multilevel models** – which we will learn

Two-step analyses – slopes as outcomes – an approximation to LMEs

1. We estimate the coefficient of the 'Portuguese' effect for each class separately

Two-step analyses – slopes as outcomes – an approximation to LMEs

1. We estimate the coefficient of the 'Portuguese' effect for each class separately
2. Then we take those per-class coefficients as the outcome variable

Two-step analyses – slopes as outcomes – an approximation to LMEs

1. We estimate the coefficient of the 'Portuguese' effect for each class separately
2. Then we take those per-class coefficients as the outcome variable
3. Examine if the per-class estimates of the experimental effect are reliably different from zero
4. or if the per-class estimates vary in relation to some variable

Slopes-as-outcomes analyses were (are) common but there is an obvious problem

- The standard errors vary widely

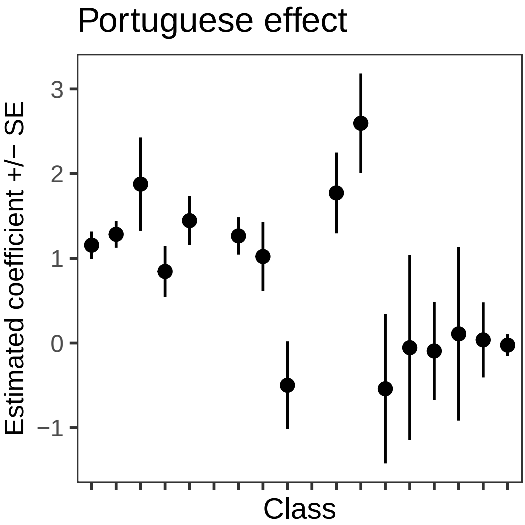

Slopes-as-outcomes analyses were (are) common but there is an obvious problem

- The standard errors vary widely
- but the two-step modelling approach takes into account the between-class differences in estimates

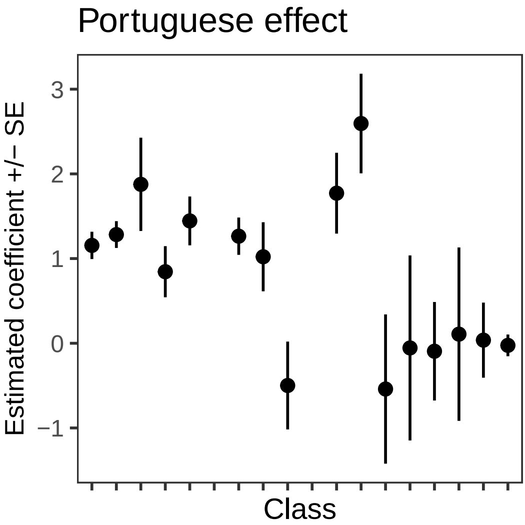

Slopes-as-outcomes analyses were (are) common but there is an obvious problem

- The standard errors vary widely
- but the two-step modelling approach takes into account the between-class differences in estimates
- but cannot account for variation in the uncertainty about those estimates – see the variation in the line intervals (SEs)

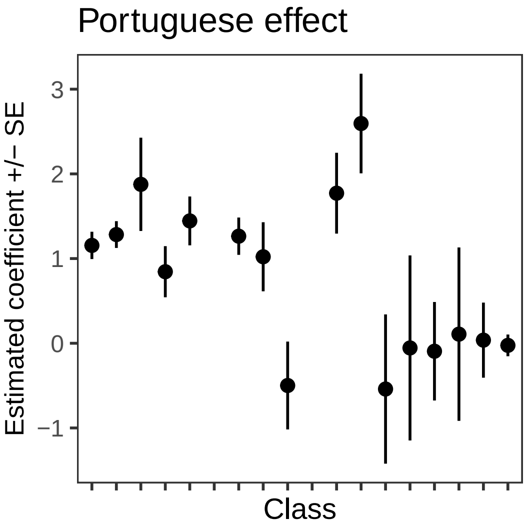

Multilevel models incorporate estimates of independent variables plus estimates of the random variation between classes in intercepts and slopes

We model the intercept as

$$\beta_{0j} = \gamma_0 + U_{0j} \tag{3}$$

- where β_{0j} values equal γ_0 the average intercept

Multilevel models incorporate estimates of independent variables plus estimates of the random variation between classes in intercepts and slopes

We model the intercept as

$$\beta_{0j} = \gamma_0 + U_{0j} \tag{3}$$

- where β_{0j} values equal γ_0 the average intercept
- plus U_{0j} the deviations between that average intercept and the intercept for each class

Multilevel models incorporate estimates of independent variables plus estimates of the random variation between classes in intercepts and slopes

We model the coefficient of the skill effect as

$$\beta_{1j} = \gamma_1 + U_{1j} \tag{4}$$

- where β_{1j} values equal γ_1 the average slope

Multilevel models incorporate estimates of independent variables plus estimates of the random variation between classes in intercepts and slopes

We model the coefficient of the skill effect as

$$\beta_{1j} = \gamma_1 + U_{1j} \tag{4}$$

- where β_{1j} values equal γ_1 the average slope
- plus U_{1j} the deviations between that average slope and the slope for each class

These models can be combined

$$y_{ij} = \gamma_0 + \gamma_1 X_{ij} + U_{0j} + U_{1j} X_{ij} + e_{ij} \quad (5)$$

- So: the English grade observed for each child y_{ij} can be predicted given the intercept γ_0

These models can be combined

$$y_{ij} = \gamma_0 + \gamma_1 X_{ij} + U_{0j} + U_{1j} X_{ij} + e_{ij} \quad (5)$$

- So: the English grade observed for each child y_{ij} can be predicted given the intercept γ_0
- plus the average relationship between English and Portuguese language skills γ_1

These models can be combined

$$y_{ij} = \gamma_0 + \gamma_1 X_{ij} + U_{0j} + U_{1j} X_{ij} + e_{ij} \quad (6)$$

- plus adjustments to capture the difference between the average grade overall and the average grade for the child's class U_{0j}

These models can be combined

$$y_{ij} = \gamma_0 + \gamma_1 X_{ij} + U_{0j} + U_{1j} X_{ij} + e_{ij} \tag{6}$$

- plus adjustments to capture the difference between the average grade overall and the average grade for the child's class U_{0j}
- plus the difference between the average slope of the Portuguese effect and the slope of that effect for their class $U_{1j}X_{ij}$

These models can be combined

$$y_{ij} = \gamma_0 + \gamma_1 X_{ij} + U_{0j} + U_{1j} X_{ij} + e_{ij} \quad (6)$$

- plus adjustments to capture the difference between the average grade overall and the average grade for the child's class U_{0j}
- plus the difference between the average slope of the Portuguese effect and the slope of that effect for their class $U_{1j}X_{ij}$
- plus any residual differences between the observed and the model predicted English grade e_{ij}

The linear mixed-effects model in R – the **lmer** function

```
porto.lmer1 <- lmer(english ~ portuguese +
                              (portuguese + 1|class_number),
                              data = BAFACALO_DATASET)
summary(porto.lmer1)
```

- porto.lmer1 <- lmer(..) estimate a linear mixed-effects model

The linear mixed-effects model in R – the **lmer** function

```
porto.lmer1 <- lmer(english ~ portuguese +
                             (portuguese + 1|class_number),
                             data = BAFACALO_DATASET)
summary(porto.lmer1)
```

- porto.lmer1 <- lmer(..) estimate a linear mixed-effects model
- *english ∼ portuguese* estimate *fixed effect* – the effect of portuguese scores on english scores

The linear mixed-effects model in R – the **lmer** function

```
porto.lmer1 <- lmer(english ~ portuguese +
                              (portuguese + 1|class_number),
                              data = BAFACALO_DATASET)
summary(porto.lmer1)
```

- (...|class_number) estimate random differences between sample groups (classes) coded by the class number variable

The linear mixed-effects model in R – the **lmer** function

```
porto.lmer1 <- lmer(english ~ portuguese +
                              (portuguese + 1|class_number),
                              data = BAFACALO_DATASET)
summary(porto.lmer1)
```

- (...|class_number) estimate random differences between sample groups (classes) coded by the class number variable
- 1 |class_number) estimate random differences between sample groups (classes) in intercepts coded 1

The linear mixed-effects model in R – the **lmer** function

```
porto.lmer1 <- lmer(english ~ portuguese +
                            (portuguese + 1|class_number),
                            data = BAFACALO_DATASET)
summary(porto.lmer1)
```

- (...|class_number) estimate random differences between sample groups (classes) coded by the class number variable
- 1 |class_number) estimate random differences between sample groups (classes) in intercepts coded 1
- (portuguese... |class_number) estimate random differences between sample groups (classes) in slopes of the portuguese effect

The linear mixed-effects model in R – the **lmer** function

- The summary shows coefficient estimates like in a linear model summary *but no p-values*

```
## Linear mixed model fit by REML ['lmerMod']
## Formula: english ~ portuguese + (portuguese + 1 | class_number)
##    Data: BAFACALO_DATASET
##
## REML criterion at convergence: 2104.3
##
## Scaled residuals:
##      Min       1Q   Median       3Q      Max
## -2.81321 -0.59584  0.04359  0.60018  2.23722
##
## Random effects:
##  Groups       Name        Variance Std.Dev. Corr
##  class_number (Intercept) 341.4803 18.479
##               portuguese    0.3295  0.574   -0.98
##  Residual                 493.1009 22.206
## Number of obs: 230, groups:  class_number, 18
##
## Fixed effects:
##             Estimate Std. Error t value
## (Intercept)  25.2837     6.2669   4.034
## portuguese    0.6590     0.1729   3.811
##
## Correlation of Fixed Effects:
##            (Intr)
## portuguese -0.943
```

The linear mixed-effects model in R – the **lmer** function

- The summary shows coefficient estimates like in a linear model summary *but no p-values*
- Plus error variance – difference for each class between the average and the class intercept

```
## Linear mixed model fit by REML ['lmerMod']
## Formula: english ~ portuguese + (portuguese + 1 | class_number)
##    Data: BAFACALO_DATASET
##
## REML criterion at convergence: 2104.3
##
## Scaled residuals:
##      Min       1Q   Median       3Q      Max
## -2.81321 -0.59584  0.04359  0.60018  2.23722
##
## Random effects:
##  Groups       Name        Variance Std.Dev. Corr
##  class_number (Intercept) 341.4803 18.479
##               portuguese    0.3295  0.574   -0.98
##  Residual                 493.1009 22.206
## Number of obs: 230, groups:  class_number, 18
##
## Fixed effects:
##             Estimate Std. Error t value
## (Intercept)  25.2837     6.2669   4.034
## portuguese    0.6590     0.1729   3.811
##
## Correlation of Fixed Effects:
##            (Intr)
## portuguese -0.943
```

The linear mixed-effects model in R – the **lmer** function

- The summary shows coefficient estimates like in a linear model summary *but no p-values*
- Plus error variance – difference for each class between the average and the class intercept
- Plus error variance – difference for each class between the average and the class slope

```
## Linear mixed model fit by REML ['lmerMod']
## Formula: english ~ portuguese + (portuguese + 1 | class_number)
##    Data: BAFACALO_DATASET
##
## REML criterion at convergence: 2104.3
##
## Scaled residuals:
##      Min       1Q   Median       3Q      Max
## -2.81321 -0.59584  0.04359  0.60018  2.23722
##
## Random effects:
##  Groups       Name        Variance Std.Dev. Corr
##  class_number (Intercept) 341.4803 18.479
##               portuguese    0.3295  0.574   -0.98
##  Residual                 493.1009 22.206
## Number of obs: 230, groups:  class_number, 18
##
## Fixed effects:
##             Estimate Std. Error t value
## (Intercept)  25.2837     6.2669   4.034
## portuguese    0.6590     0.1729   3.811
##
## Correlation of Fixed Effects:
##            (Intr)
## portuguese -0.943
```

Conclusions – Linear Mixed-effects models: Why learn about them?

1. Psychological studies frequently result in hierarchically structured data
2. Structure can be understood in terms of the grouping of observations
    - as when there are multiple observations per group, participant or stimulus
3. Multilevel or mixed-effects models can be specified by the researcher
    - to include random effects parameters that capture unexplained differences between participants or other sampling units
    - in the intercepts or the slopes of explanatory variables

Summary – Week 17: getting started

- *What, why, when*
  - *What, why* – What are multilevel models? – Why do we use them?
  - *When* should they be applied?
- *How* do we conduct multilevel models? – we get started this week, develop skills over next weeks
- Next week, we will tackle the ideas in a bit more depth