

Workshop 5: Linear mixed-effects models

In this workshop we will fit linear mixed-effects models to data in R. Linear mixed-effects models are used when you have random effects, which occurs when multiple measurements are made on randomly sampled units. Thus the measurements are grouped rather than obtained by independent random samples. The units or groups are assumed to be randomly sampled from a “population of groups”. Example situations include

- when you divide up plots and apply separate treatments to the parts (plot is the random group)
- when your sampling design is nested, such as quadrats within transects; transects within woodlots; woodlots within districts (transects, woodlots, and districts are all random groups)
- when you take measurements on related individuals (family is the random group)
- when you measure subjects repeatedly (subject is the random group)

Linear models for mixed effects are implemented via the [lme4](#)/[lmerTest](#) package. An alternative option is to use the [lme](#) method in the [nmlme](#) package, but the approximations used to calculate F -statistics and degrees of freedom in [lme4](#) are slightly more accurate, especially when sample size is not large.

Repeatability of a sexual signal trait

This first data set was extracted from a paper by Griffith and Sheldon (2001, *Animal Behaviour* 61: 987–993), who measured the white forehead patch of 30 male collared flycatchers in two years on the Swedish island of Gotland. The patch is important in mate attraction, but varies in size from year to year. Our goal here will be to estimate the repeatability of patch length (mm). The data are [here](#).

Read and examine the data

1. Read the data from the file.
2. View the first few lines of data to make sure it was read correctly.
3. Use the [interaction.plot](#) command to create a plot showing the pair of measurements for each individual flycatcher in the two years of study. Is there evidence of measurement variability between years?

Fit a linear mixed-effects model

1. Fit a linear mixed-effects model to the data, treating the individual birds as the random groups.

Run it with and without the fixed effect of year. To include year as a fixed effect, make sure to convert year to a factor in R so it is not treated as a numeric variable. Recalculate repeatability with this model as described in steps (2) and (3) below. How is the interpretation of repeatability changed? Extract parameter estimates from the saved `lmer` object (the commands are the same ones we used with `lm`). Inspect the output for the random effects and for the fixed effect.

2. In the output, examine the standard deviations for the random effects. There should be two standard deviations: one for “(Intercept)” and one for “Residual”. This is because the mixed effects model has two sources of random variation: variation among repeat measurements within birds, and true variation among birds in their patch lengths. Which of these two sources corresponds to “(Intercept)” and which to “Residual”?
3. Also examine the output for the fixed effect results. The only fixed effect in the model formula is the grand mean of all the patch length measurements. It is called “(Intercept)”, but don’t confuse with the intercept for the random effects. The fixed effect output gives you the estimate of the grand mean and a standard error for that estimate. Notice how the fixed effect output provides estimates of means, whereas the random effects output provides estimates of variances (or standard deviations).
4. Extract the variance components from the fitted model and estimate the repeatability of patch length from year to year. Hint: use the `VarCorr` command to extract the variance components.
5. Interpret the measure of repeatability obtained in the previous step. If the repeatability you obtained is less than 1.0, what is the source of the variation among measurements within individuals. Is it measurement error alone?
6. Produce a plot of residuals against fitted values. Notice anything odd? There seems to be a slightly positive trend. This isn’t a mistake, but results from “shrinkage” of the best linear unbiased predictors (BLUPs). Consult the lecture notes for additional information on what is happening.

Goldie’s vision

Cronly-Dillon and Muntz (1965; *J. Exp. Biol.* 42: 481-493) used the optomotor response to measure color vision in the goldfish. We looked at some of these data already in lecture. Here we will fit a model to the data and include the full set of wavelengths tested. Each of 5 fish was tested at all the wavelength in random order. A large value of sensitivity indicates that the fish can detect a low light intensity. An important feature of the optomotor response is that fish don’t habituate, and it is unlikely that a measurement of visual sensitivity under

one wavelength would have an effect on later measurements at another wavelength. The data are [here](#).

Read and examine the data

1. Read the data from the file, and view the first few lines to make sure it was read correctly.
2. Use the `interaction.plot` command to compare the responses of individual fish across the different light wavelength treatments.
3. What type of experimental design was used? This will determine the linear mixed model to use when fitting the data.

Fit a linear mixed-effects model

1. Fit a linear mixed-effects model to the data.
2. Use the `visreg` command (you'll need to load the `visreg` library first) to plot the fitted values (with wavelength on the x -axis). Note that the predicted lines for different fish are parallel (because you have fit a model without an interaction term - including an interaction term is not possible with these data because there is only one measurement per fish per wavelength) and also on top of one another (because there is almost no estimated variance for the 'fish' random effect).
3. Use the `fitted(out)` command to compare the predicted values from (2) with the observed data. The difference between the predicted and observed values should be equal to the residuals (which you can access with `resid(out)`).
4. What assumptions are you making in (1)? Create a plot of residuals against fitted values to check one of these assumptions.
5. Extract parameter estimates for the fixed effect (wavelength) from the saved `lmer` object using the `fixef` command. The coefficients have the same interpretation as in the case of a categorical variable analyzed using `lm` (arbitrarily, the group "nm426" is set as the "control").
6. Inspect the output for the random effects using the `ranef(out)` command (this should give you the individual fish effects). In our mixed model, we have two sources of random error. What are they? Which of them corresponds to the "(Intercept)" and which to the "Residual" in the output? Notice that the estimated standard deviation for one of the sources of variation is very small.

Yukon yarrow

The Kluane project experimentally investigated the effects of fertilization and herbivory on vegetation dynamics in the boreal forest ecosystem of Kluane National Park in the Yukon

(Krebs, C.J., Boutin, S. & Boonstra, R., eds (2001a) Ecosystem dynamics of the Boreal Forest. The Kluane Project. Oxford University Press, New York). The current data are from a study of the effects of plant resources and herbivory on the defensive chemistry of understory plant species.

Each of sixteen $5 \times 5m$ plots was randomly assigned one of four treatments: 1) surrounded by a fence enclosure to exclude herbivores; 2) fertilized with N-P-K fertilizer; 3) fenced and fertilized; and 4) untreated control. Each of the 16 plots was then divided in two. One side of each plot (randomly chosen) received the treatment continually over the 20 years of study. The other half of each plot received the treatment for the first ten years, after which it was left to revert to its untreated state. The data to be analyzed here record the concentration of phenolics (a crude measure of plant defense compounds) in yarrow (*Achillea millefolium*), a herb common in the plots. The measurement units are mg Tannic Acid Equivalents per g dry weight. The data were kindly provided by P. de Koning and R. Turkington. The data are [here](#).

Visualize the data

1. Read the data from the file.
2. Inspect the first few lines of data. Plot and treatment are self-explanatory. Treatment is given as a single variable with four levels (let's stick with this approach rather than model as two variables, enclosure and fertilizer, with a 2×2 factorial design). Duration indicates whether the half-plots received the treatment for the full 20 years or whether the treatment was stopped ("reversed") after 10 years. The variable "phen.ach" is the concentration of phenolics in yarrow.
3. What type of experimental design was used? This will determine the linear mixed model to use when fitting the data.
4. Draw a graph to illustrate the concentrations of phenolics in yarrow in the two duration categories for each treatment. There aren't many data points in each combination of treatment and duration levels, so a grouped strip chart is probably a better choice than a grouped box plot.

Fit a linear mixed-effects model

1. Fit a linear mixed model to the data without an interaction between treatment and duration. Use the log of phenolics as the response variable, as the log-transformation improved the fit of the data to linear model assumptions. Visualize the model fit to the data.
2. Now repeat the model fitting, but this time include the interaction between treatment and duration. Visualize the model fit to the data. What is the most noticeable difference between the two model fits, one with the interaction and the other without? Describe what including the interaction term "allows" that the model without an interaction term does not. Which model appears to fit the data best?

3. Check the residuals in relation to the fitted values to investigate a key assumption of linear mixed models.
4. Extract the parameters from the fitted model object. In the output generated, you will see two quantities given for “Std.Dev” under the label “Random effects”. Explain what these quantities refer to.
5. Use the `drop1` command (with option `test='Chisq'`). Because we have random effects, we can no longer use an F -test to test significance. If you're feeling ambitious, try fitting the model using the `lme` command from the `nlme` package. Do your results agree?