

Use of Linear Mixed Models for Experimental Design

What are linear mixed models?

They are a group of linear models that include:

- One dependent variable, that is continuous (usually labeled as Y or y in textbooks)
- fixed components
 - continuous variables, and/or class variables represented by dummy (indicator) variables;
 - fixed-effects in experimental design, predictor variables in regression, usually labeled as X or x ;
 - associated coefficients are labeled as β in most texts.
- error term
 - usually labeled as ε (use e if this is estimated errors, not population errors)
 - covariance matrix: variances and covariances of the errors; labeled the R matrix in many mixed models text books
 - error terms follow a normal distribution
 - error terms may have unequal variance, and /or correlations (time and/or space) between error terms
 - error terms are a random component.

and may include, also:

- random components
 - covariance matrix (variances and covariances of these random components) is labeled the G matrix in many texts
 - the “variables” (really a design matrix) are labeled as Z , with associated coefficients “ u ”.
 - these also follow a normal distribution
 - some models have only random components, and no fixed components

Aside: In math symbols, this becomes:

$$\mathbf{y} = \boldsymbol{\beta}\mathbf{x} + \mathbf{u}\mathbf{Z} + \boldsymbol{\varepsilon} \qquad \mathbf{V}(\mathbf{y}) = \mathbf{G}'\mathbf{Z}\mathbf{G} + \mathbf{R}$$

- Estimates of all parameters:
 - the fixed component coefficients (including the intercept),
 - the variances for:
 - the random components variances and covariances; and random-effects coefficients
 - variances (and covariances) of the error term
- are estimated using maximum likelihood

Likelihood

Given a set of the estimated parameters (coefficients and variances/covariances), what is the chance that we would get the data that we did get?

For a discrete distribution of y (not the case in linear mixed models), this would be a probability for the first observation X the prob of the second observation, etc. to the last observation – between 0 and 1.

For a continuous distribution, e.g., normal, this is the value of the probability density function for the first observation X the probability density function for the second observation, etc to the last observation – not necessarily less than 1.

Maximum Likelihood

Change the set of estimated parameters until we get the largest possible likelihood.

Often easier to take the logarithm of the likelihood to do this – most packages report the log likelihood, or -2 X log likelihood.

Searching for the Maximum Likelihood

Most packages get the maximum likelihood by:

- Searching for a set of all of these estimated parameters that will result in the maximum likelihood of obtaining the data that we did get (ML method)

OR

- Finding estimates of the fixed component coefficients first (sometimes using least squares methods), and then using the residuals from that to get the random components (REML).

Because this is a search to find a solution (the estimates that give the maximum likelihood), the search proceeds by :

- getting estimates, calculating the (log) maximum likelihood (one iteration),
- altering the estimates, and recalculating the maximum likelihood (another iteration), and
- so on, until the estimates don't change (or this may stop based on the likelihood does not change).

However, the search may not converge –

- means that the estimates are not becoming the same over the iterations of the search.
- You may need to:
 - increase the number of iterations,

- change the way the search is done (e.g., Marquardt is one method for searching that is commonly used)
- It may mean that your model is not correctly specified, or it is just very hard to find a solution if your models are very complex.

The search may converge, but with the statement that the “Hessian is not positive definite”

- This will mean that the variance estimates are not reliable.
- This can occur with a complex model, or when the model is not correctly specified.

Mixed models for experimental design

Linear mixed models enable us to get estimates for mixed-effects models, including:

- testing the fixed-effects factors for interactions, and main effects (Type III SS, F-tests). Using a mixed effects model, correct F-tests will be calculated.
- Get t-tests for pairs of means using the correct denominator Mean Squares (same as the one used in the F-test)
- Get estimates of the variances for the random effects, including the variance of the residual error.
- Testing assumptions: bit harder to do!
 - Use “white noise residuals” and test for equal variance, normality, and independence of residuals (over time or space if you suspect this)
 - Check the log likelihood – should be better (higher log likelihood OR lower -2 log L) as you better meet the assumptions.

Examples of Experiments with Mixed Effects

Example 1: Completely randomized design with random or mixed effects

- Factor A, (three levels of fertilization: A1, A2, and A3) (J=3) – fixed-effects.
- Factor B (four species: B1, B2, B3 and B4) (K=4) Random-effects. Crossed: 12 treatments
- Four replications per treatment (n=4) for a total of 48 experimental units
- Measured Responses: height growth in mm

species is random -- these are a few of the species that we are interested in and we wish to look at the variance in height growth that is due to species.

fertilization is fixed – these are the only levels of interest and we wish to compare these specific levels of fertilizers in terms of height growth.

Example 2: Randomized Block Design with replicates in each block

- Factor A (three types of food: A1 to A3)
- Two labs on which the experiment is performed (blocks).
- Randomization of Factor A is restricted to within labs.

Lab 1

Lab 2

A1 = 6	A1=5		A3=11	A3=12
A3=10	A2=8		A1=4	A2=9
A2=7	A3=12		A2=8	A1=5

Response variable: weight gain of fish (kg)

Experimental unit: one tank of fish; 6 tanks in each lab

Blocks are random as we are not interested in comparing the two labs, just in estimating the variance due to different labs.

The type of foods is a fixed effect as we are interested in comparing these specific foods.

Example 3: CRD: One Factor Experiment, Fixed Effects with subsampling

Example from Kutner et al., Applied linear models, textbook:

- Have three temperatures: low, medium, and high (fixed effect)
- For each, we have two experimental units (batches) (random effect)
- For each batch, we have three loaves of bread (random effect)
- The response variable is crustiness of bread.

Data:

temp	batch	observation	y _{ijl}
low	1	1	4
low	1	2	7
low	1	3	5
low	2	1	12
low	2	2	8
low	2	3	10
medium	1	1	14
medium	1	2	13
medium	1	3	11
medium	2	1	9
medium	2	2	10
medium	2	3	12
high	1	1	14
high	1	2	17
high	1	3	15
high	2	1	16
high	2	2	19
high	2	3	18

There are two levels in a hierarchy: batches and then loaves for each batch.

The treatment is applied to the batch, so the batch is the experimental unit.

We could average the loaves for each batch, and then analysis this as a very simple one factor completely randomized design OR use a mixed linear model and separate out the two levels, batch and then loaf within batch.

References:

Littell, R.C., G. A. Milliken, W.W. Stroup, and R.D. Wolfinger. 1996. SAS system for Mixed Models. SAS Institute Inc., Cary, NC. [does not have R code, but has excellent examples of using mixed models for experiments]
Pineiro, J.C. and D.M. Bates. 2000. Mixed-effects models in S and S-plus. Springer, New York. [has descriptions of the examples we will use in the workshop]
Schabenberger, O. and F. J. Pierce. 2002. Contemporary Statistical Models. CRC Press, New York, Chapter 7. "Linear mixed models for clustered data."