

A short course in Longitudinal Data Analysis

ESRC Research Methods and Short Course

Lab 3 - June 2011

The nlme library in R

Load the package nlme.

```
> library(nlme)
```

We can fit a general linear mixed effects model using the `lme` function. The following example fits a linear time-trend with a random intercept for the random effect component of the model, using the method of maximum likelihood.

```
> model1 <- lme(Y ~ time, data = mental1, random = ~1 | id, method="ML", na.action=na.omit)
summary(model1)
```

Linear mixed-effects model fit by maximum likelihood

Data: mental1

	AIC	BIC	logLik
	1533.086	1546.398	-762.5432

Random effects:

Formula: ~1 | id

(Intercept) Residual

StdDev: 9.815321 7.720669

Fixed effects: Y ~ time

	Value	Std.Error	DF	t-value	p-value
(Intercept)	55.50867	1.5943630	155	34.81558	0.0000
time	0.35609	0.2401043	155	1.48305	0.1401

Correlation:

(Intr)

time -0.314

Standardized Within-Group Residuals:

Min	Q1	Med	Q3	Max
-3.11678026	-0.53936453	-0.05250011	0.53014796	3.13757774

Number of Observations: 206

Number of Groups: 50

The `lme()` function does not use the same parameterisation for the random effects as was used in the lecture notes. To obtain the standard parameterisation the transformations are the following.

$$\begin{aligned}\nu &= \text{intercept} \\ \sigma^2 &= (\text{residual})^2 * (1 - \text{nugget}) \\ \phi &= \text{range} \\ \tau^2 &= (\text{residual})^2 * \text{nugget}\end{aligned}$$

Exercise. Fit the same model for the mean, but using the `lm()` function, which assumes that repeated measurements on the same subject are uncorrelated. Compare the standard errors returned by this method with those returned by the `lme()` function.

The following example fits a model specifying both a random intercept and random slope.

```
> model2 <- lme(Y ~ 1 + time, data = mental1, random = ~time | id,
  method = "ML",na.action=na.omit)
```

Exercise. Does the random intercept and slope model fit the `mental1` data better than the random intercept model?

Exercise. Fit a model to the `mental1` data that includes a quadratic time-trend.

Exercise. Fit a suitable random effects model to the complete `mental` data-set. Use your model to compare the three treatments and to decide which, if any, of the treatments achieve the clinical objective of reducing the mean response at week 8 to at most 80% of the mean response at week 0.

To exit from R, type `q()`, but **remember to save your workspace** for the later practical classes.

PJD, 25.05.2011