

Mixed effects models

The basic theory and application in R

Mitchel van Loon

Research Paper Business Analytics

Mixed effects models

The basic theory and application in R

Author:
Mitchel van Loon

Research Paper Business Analytics

Supervisor:
Eduard Belitser

Vrije Universiteit Amsterdam
Faculteit der Exacte Wetenschappen
Studierichting Business Analytics
De Boelelaan 1081a
1081 HV Amsterdam

October 2014

Preface

This paper serves as a Research Paper for the master Business Analytics at VU University Amsterdam. The goal of this paper is to get acquainted with both the theory of mixed effects models and the tools in R for mixed effects models. Mixed effects models can be applied when modelling factors that are generally not of importance for the overall research. I would like to thank my supervisor Eduard Belitser for his guidance throughout the process of writing this paper.

Summary

Statistical models are sometimes applied to data that cope with individual deviation of observations. These deviations are generally not important but they are present. Thus, one should keep these effects in mind. The goal of this research is to get a general view of the theory and application of statistical models that are able to cope with these individual differences, the mixed effects model.

The mixed effects model is similar to the often applied regression model, except for one main addition, the random effects part. This is the part that explains the individual effects of the different observations. Just as the fixed part this can be split into coefficients and variables. However, the coefficients are random and assumed to be multivariate normally distributed. Hence, they are similar to the error term in linear regression.

The estimation of the parameters of mixed effects models is either done by the maximum likelihood estimator or by the restricted maximum likelihood estimator, which is the default in R. To perform analysis on mixed effects models in R, the `lmer` function from the `lme4` library can be applied.

Index

Preface.....	3
Summary	4
Introduction.....	6
1 The linear regression model.....	6
2 The linear mixed effects model.....	7
3 Sleep study example.....	11
4 Muscle building example.....	14
Discussion	19
References.....	20

Introduction

Regression models are powerful statistical tools to analyze data and use this data for predictions. The linear regression model is often applied here. However, this model may not always capture the data sufficiently enough. It is often not possible to fully control certain circumstances of data capturing so there may be variation among measurements. This part of a model is usually accounted for by the error term in a linear regression model. However, some of this variation can be explained in more detail and a mixed effects model may be appropriate for the given data. Several questions are to be asked here, such as when to apply mixed effects models and how should this be approached.

This paper will discuss the differences between the standard linear regression model and the mixed effect model. This will be done in both theory and application. For instance, in a medical trial patients may react differently on a certain drug. So one could ignore this random effect and part of this randomness would probably end up in the error term of a regression model. However, one could also try to capture this random effect more precisely via a mixed effects model.

In chapter 1 we will discuss the basic regression model. Chapter 2 will discuss the mixed effects model compared to the basic regression model. In chapter 3 and 4 two examples of the application of mixed effects models will be worked out with the statistical program R. These examples will be done based on data from Patterns of performance degradation and restoration during sleep restriction and subsequent recovery: a sleep dose-response study by Belenky et al. (2003) and Applied Longitudinal Analysis by Fitzmaurice et al. (2004).

1 The linear regression model

The linear regression model is the most common form of statistical modelling and is applied in various fields including economics, healthcare and science like social and behavioral sciences. The model is applied when there is a linear relationship between the independent variables and a dependent variable (observation).

The basic linear regression model has the following form for each observation i :

$$Y_i = \beta_1 X_{i,1} + \dots + \beta_p X_{i,p} + \epsilon_i$$

Here β_1, \dots, β_p are the unknown parameters, $X_{i,j}$'s are some known design variables, Y_i 's are the observations and ϵ_i 's are the error terms, where $\epsilon_1, \dots, \epsilon_n \sim N(0, \sigma^2)$ and independent. Rewriting this in matrix notation gives the following:

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon},$$

where $\mathbf{y} = (Y_1, \dots, Y_n)^T$; $\mathbf{X} = (X_{ij})$; $i = 1, \dots, n$; $j = 1, \dots, p$; $\boldsymbol{\beta} = (\beta_1, \dots, \beta_p)^T$; and $\boldsymbol{\epsilon} = (\epsilon_1, \dots, \epsilon_n)^T$.

In order to set up a model we need to estimate the coefficients. This can be done by minimizing the sum of squared errors. We then obtain the least square estimator of the coefficients, commonly given by:

$$\hat{\beta} = (X^T X)^{-1} X^T y.$$

A condition here is that X is of the full column rank, thus $\text{rank}(X) = p$.

Another way of estimating the coefficients is the maximum likelihood estimate. This is the most commonly used estimation method. However, when ϵ_i are normally distributed, the resulting estimators of the maximum likelihood estimation are identical to the least squares estimators.

If a model is set up, one might want to test the fit of the model and other aspects such as the coefficients. An F-test can be used to check the overall fit of the model and a t-test can be used to check coefficients individually. Most statistical programs such as R include these tests in the results of a fitted model.

2 The linear mixed effects model

The linear mixed model would have the following form for each observation i :

$$Y_i = \beta_1 X_{i,1} + \dots + \beta_p X_{i,p} + b_1 Z_{i,1} + \dots + b_q Z_{i,q} + \epsilon_i$$

Here $\beta_1 \dots \beta_p$ are the unknown parameters of the fixed effects, $X_{i,j}$'s are the known design variables and Y_i 's are the observations. Additionally $b_{i,j}$'s are the random coefficients and $Z_{i,j}$ are the random effect design variables. Finally, the ϵ_i 's are the error terms, where $\epsilon_1, \dots, \epsilon_n \sim N(0, \sigma^2)$ and independent. Fox (2002) states that the random coefficients should be viewed similar to the error terms rather than the fixed coefficients. Hence, there is a variation component in addition to the error term seen in traditional regression models such as the linear regression model. We assume $b_{i,j} \sim N(0, \psi^2)$ where ψ^2 resembles the variances. When we rewrite all assumptions in matrix notation we get the following setup:

$$Y = X\beta + Zb + e,$$

X is the $n \times p$ design matrix,
 β is the $p \times 1$ parameter vector,
Z is the $n \times q$ design matrix,
 $b \sim N_q(0, \sigma^2 D),$
D is the $q \times q$ parameter matrix,
 σ^2 is the parameter,
 $e \sim N_n(0, \sigma^2 I),$
b, e independent.

The mixed effects model mixes random effects with the fixed effects of a standard regression model. Winter (2013) explains random effects as something that is usually nonsystematic and unpredictable, hence random influence on the data. Fixed effects however, are predictable and systematic.

Mixed effects models may be useful when analyzing data of certain experiments, depending on the way these experiments are designed. When there are factors that you are not interested in per se affecting an observation, this is called a blocked design. When for instance, you would like to test a certain medicine in a medical trial. A random effect is then

the individual response of a test subject. This is usually tested over a certain period of time, i.e. longitudinal data. There may be certain groups of subjects that perform differently when the same experiments are done on different groups, i.e. repeated measures. Another example is nested or crossed data.

Blocked design

When analyzing blocked data there are variables that are not of interest to the experiment. Say we want to measure the effects of variable X1 on the variable Y, however we also measure another variable X2. This variable may be of influence on our measurement Y. Thus, we would like to use the measured variable X2 to improve our knowledge. This is where mixed effects models come in. The blocked design assumes that the X2 measurements are outside the control of the experiment, and thus we assume that they are random. Our model would then have a fixed effects variable X1 and a random effects variable X2.

Nested design

When analyzing nested data there are variables nested in other variables. For instance, when there are 2 people performing the exact same experiment they may still provide different results. One would then say the experiments are nested in the person.

Repeated measures design

In the repeated measurements model the measurements are grouped while the same experiment is performed for different groups. Measurements within a group are dependent and measurements between groups are independent. This can be modelled by a random effect for each group, independently across groups. Thus, the random effect would explain the differences between the different groups.

Longitudinal design

In longitudinal analysis groups correspond to measurements on the same individual over time. Random effects permit individuals to follow different response trajectories. These differences may be caused by variation of measurement circumstances, such as the number and timing of observations for individuals.

Based on the different types of data we may have different structures of parameter matrix D. The basic structure is a diagonal matrix. This is the case when each of the random effect parameters are independent. The matrix $\sigma^2 D$ then contains the variance of each random parameter at the diagonal. For a mixed effects model with q independent random effect parameters the matrix is of the following form:

$$\sigma^2 D = \begin{bmatrix} d_1 \sigma^2 & 0 & \dots & 0 \\ 0 & d_2 \sigma^2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & d_q \sigma^2 \end{bmatrix}$$

Similar to the linear regression model we need to obtain certain estimates. For mixed effects models the most common estimates are the maximum likelihood (ML) estimate and the restricted maximum likelihood (REML) estimate, the latter being the default estimate in R.

Maximum likelihood estimates are obtained by maximizing a likelihood function. Given the parameters in the mixed effects model the likelihood can be constructed subject to the multivariate normality assumption of the dependent variable Y_i . The log likelihood function is the log density function of the normal distribution $N(X\beta, \sigma^2(ZDZ^T + I))$ viewed as a function of the parameters β , σ^2 and D . This function is given by:

$$-\frac{n}{2}\log(2\pi) - \frac{1}{2}\log(\det(\sigma^2(ZDZ^T + I))) - \frac{1}{2\sigma^2}(Y - X\beta)^T(ZDZ^T + I)^{-1}(Y - X\beta)$$

To derive this function we start at the assumption that we have for the ML estimate:

$$Y \sim N(X\beta, \sigma^2(ZDZ^T + I))$$

The first thing we see is that the mean equals the fixed effects part of $Y = X\beta + Zb + e$.

$$EY = X\beta$$

The covariance is obtained using the independency of b and e as follows:

$$Cov(Y) = Cov(Zb) + cov(\epsilon) = ZCov(b)Z^T + Cov(\epsilon) = \sigma^2(ZDZ^T + I)$$

Now all that is left is to apply the results in the likelihood function of the multivariate normal distribution and to take the log to obtain the final formula. The likelihood function is as follows:

$$X \sim \mathcal{N}(\mu, \Sigma) \Leftrightarrow X \sim f(x) = (2\pi)^{-\frac{n}{2}}(\det(\Sigma))^{-\frac{1}{2}}\exp\left\{-\frac{1}{2}(x - \mu)^T\Sigma^{-1}(x - \mu)\right\}$$

Having the general likelihood function of a multivariate normal distribution we can substitute the parameters for the ML estimation as follows:

$$Y \sim \mathcal{N}(X\beta, \sigma^2(ZDZ^T + I)) \Leftrightarrow Y \sim f(y) = (2\pi)^{-\frac{n}{2}}(\det(\sigma^2(ZDZ^T + I)))^{-\frac{1}{2}}\exp\left\{-\frac{1}{2}(y - X\beta)^T(\sigma^2(ZDZ^T + I))^{-1}(y - X\beta)\right\}$$

Taking the log of this result we obtain the ML log-likelihood function given.

Another way of estimating the parameters is the REML estimate. West et al. (2007) state that the REML estimate is usually preferred to the ML estimate: “, because it produces unbiased estimates of covariance parameters by taking into account the loss of degrees of freedom that results from estimating the fixed effects in β .” Another point here is that the approach based on REML makes the likelihood function independent of β . Hence, the REML

estimates are based on optimization of the following REML log-likelihood function of the parameters σ^2 and D :

$$-\frac{n}{2}\log(2\pi) - \frac{1}{2}\log(\det(\sigma^2 K(ZDZ^T + I)K^T)) - \frac{1}{2\sigma^2}(KY)^T(K(ZDZ^T + I)K^T)^{-1}(KY)$$

Where,

$$KY \sim N_k(0, \sigma^2 K(ZDZ^T + I)K^T)$$

given a matrix K such that $KX = 0$.

To derive the REML log-likelihood function we multiply our model function $Y = X\beta + Zb + e$ with the matrix K and we apply that $KX = 0$ to obtain the following:

$$KY = KX\beta + kZb + K\epsilon = 0 + KZb + K\epsilon$$

Hence we obtain a mean of 0 and the covariance becomes the following:

$$\text{Cov}(KY) = K\text{Cov}(Y)K^T = \sigma^2 K(ZDZ^T + I)K^T$$

Resulting in the assumption given for the REML estimate.

$$KY \sim N_k(0, \sigma^2 K(ZDZ^T + I)K^T)$$

Now all that is left is to apply the results in the likelihood function of the multivariate normal distribution and to take the log to obtain the final formula. Having the general likelihood function of a multivariate normal distribution we can substitute the parameters for the REML estimation as follows:

$$f(KY) = (2\pi)^{-\frac{n}{2}} (\det(\sigma^2 K(ZDZ^T + I)K^T))^{-\frac{1}{2}} \exp\left\{-\frac{1}{2}(KY)^T((\sigma^2(K(ZDZ^T + I)K^T))^{-1}(KY))\right\}$$

Taking the log of this result we obtain the REML log-likelihood function.

Because the REML function is independent of β , there is no information on the fixed effects. Thus, the generalized least square estimates are used to estimate β similarly to the estimation of β using the ML method.

A condition for both the ML and the REML method is again that X is of full rank, thus $\text{rank}(X) = p$, in which case $\hat{\beta}$, being the estimated β , exists. However, X may not be of full rank in general. If this is the case we can introduce a matrix H such that $H\hat{\beta} = 0$ and under restriction we only require $\text{rank}(X) < p$.

ML estimation is more appropriate when performing likelihood ratio tests. However, it is possible to perform such a test when REML estimation is used. Pinheiro and Bates (2000) state that this only applies if both models have been fit using REML estimation. This also

requires that the fixed effects of both models have the same specification. Hypotheses on fixed effects must always be tested comparing two models fitted by ML estimation. Another way of comparing models is a bootstrap simulation. ANOVA may be unreliable when comparing nested models. Thus, bootstrapping can be used if conclusions are in doubt. This may be more accurate, but more computational heavy as well.

Solving the ML or REML optimization problem can be done using iterative optimization algorithms such as the Newton-Raphson (NR) algorithm. Other common algorithms used in the context of mixed effects models are the Expectation-Maximization (EM) algorithm and the Fisher scoring algorithm. The statistical package R uses both the NR method and a variation on the EM method called expectation conditional maximization either (ECME).

A recent study (Li and Pourahmadi, 2012) has been performed on an alternative REML estimation (aREML) method. According to Li and Pourahmadi (2012) an advantages of the aREML is substantially reducing the parameter dimension in the estimation stage. This is done by separating the two covariance matrices and modeling them individually. Another advantage is that the aREML method requires minimal assumptions on the random effects and the structure of the covariance matrices. Li and Pourahmadi also state that the aREML method makes it possible to use databased and graphical tools to formulate models for the covariance matrices in flexible ways.

Li and Pourahmadi note that the likelihood of the aREML method is a function of the fixed effect parameters, β , and the residual error parameter in the matrix $\sigma^2 I$. They put focus on estimating σ^2 by adjusting the REML method. Next, the best linear unbiased predictors of the random effects b can be used to evaluate their distribution and formulating models for their covariance matrix D .

In the following chapters we will work out two examples of the application of mixed effects models with the statistical program R. The first example is based on data taken from Patterns of performance degradation and restoration during sleep restriction and subsequent recovery: a sleep dose-response study by Belenky et al. (2003). Bates (2010) shows a similar approach using mixed effects models on the sleepstudy data. The second example is based on data taken from Applied Longitudinal Analysis by Fitzmaurice et al. (2004). To perform analysis on mixed effects models in R, the lmer function can be called. In order to make use of the lmer function the lme4 library is needed.

3 Sleep study example

The data for the sleep study example can be found in the lme4 library and was taken from Patterns of performance degradation and restoration during sleep restriction and subsequent recovery: a sleep dose-response study by Belenky et al. (2003). The sleepstudy data consists of observations of reaction times given the amount of days that a truck driver has been restricted to only 3 hours of sleep. This experiment was done over a period of 10 days for several truck drivers.

Because each person is likely to respond differently on the amount of missed sleep this makes a good case for a mixed effects model. The random effects for each individual are the

deviations in intercept and slope of that individual compared to the general observation. The model formula would be of the following form: $\text{Reaction} = \text{intercept} + \beta_1 * \text{Days of missed sleep} + \text{individual intercept} + \text{individual effect} * \text{days of missed sleep} + \epsilon$.

Before fitting a model we can make several assumptions. For instance, one could assume that truck drivers who react fast may be less influenced by sleep deprivation than truck drivers who react slower. Thus, this indicates dependence between the individual intercept and the individual coefficient. We will fit both models to the data. Comparing these 2 models may give us an indication that, for example, truck drivers with higher initial reaction times may, on average, be more strongly affected by sleep deprivation.

To call such a mixed effects model we use the following command in R:

```
sleepmod=lmer(Reaction~1+Days+(Days | Subject),data=sleepstudy)
```

This leads to the following result:

```
Linear mixed model fit by REML ['lmerMod']  
Formula: Reaction ~ 1 + Days + (Days | Subject)  
Data: sleepstudy
```

```
REML criterion at convergence: 1743.6
```

```
Scaled residuals:
```

```
  Min   1Q   Median   3Q   Max  
-3.9536 -0.4634  0.0231  0.4634  5.1793
```

```
Random effects:
```

```
Groups   Name      Variance Std.Dev. Corr  
Subject (Intercept) 612.09  24.740  
        Days        35.07   5.922  0.07  
Residual          654.94  25.592
```

```
Number of obs: 180, groups: Subject, 18
```

```
Fixed effects:
```

```
          Estimate Std. Error t value  
(Intercept) 251.405   6.825   36.84  
Days         10.467   1.546    6.77
```

```
Correlation of Fixed Effects:
```

```
(Intr)  
Days -0.138
```

The output shows an intercept and a slope with respect to Days for both random and fixed effects. The output also shows a correlation coefficient meaning the slope of days and the intercept are dependent.

Because of the dependence of the slope and the intercept of each individual we would have a covariance matrix D that has the variances on the diagonal and the individual interaction between the intercept and the slope above and below the diagonal.

For the fixed effects the output shows that the estimates for the intercept and slope are respectively 251.405 and 10.467. This would indicate that the base reaction time is about 251 milliseconds and this increased by about 10 milliseconds for each day of sleep deprivation. For the random effects the output shows that the intercept for each truck driver may vary with a standard deviation of about 25 milliseconds, while the variation in the slope has a standard deviation of about 6 milliseconds. The output also shows a low dependence between the intercept and the slope with a correlation coefficient of 0.07.

Now if we assume that the effect of sleep deprivation is independent of the intercept of an individual we add (1 | Subject) and we add a -1 to Days:

sleepmod1=lmer(Reaction~1+Days+(1|Subject)+(Days-1|Subject),data=sleepstudy)

This means that we do have an intercept for each individual. However, this is not dependent on the actual effect of sleep deprivation (due to the -1). We obtain the following result:

Linear mixed model fit by REML ['lmerMod']

Formula: Reaction ~ 1 + Days + (1 | Subject) + (Days - 1 | Subject)

Data: sleepstudy

REML criterion at convergence: 1743.7

Scaled residuals:

Min	1Q	Median	3Q	Max
-3.9626	-0.4625	0.0204	0.4653	5.1860

Random effects:

Groups	Name	Variance	Std.Dev.
Subject	(Intercept)	627.57	25.051
Subject.1	Days	35.86	5.988
Residual		653.58	25.565

Number of obs: 180, groups: Subject, 18

Fixed effects:

	Estimate	Std. Error	t value
(Intercept)	251.405	6.885	36.51
Days	10.467	1.560	6.71

Correlation of Fixed Effects:

(Intr)

Days -0.184

The output shows the same estimated parameters for the fixed effect intercept and slope, they are still 251.405 and 10.467 respectively. For the random effects the output shows slightly different standard deviations for both the intercept and the slope. However, there is no correlation coefficient given.

The output shows that R does not give a correlation coefficient using this model due to the independence. Thus, the matrix D is a diagonal matrix containing only the variances. We can now compare the 2 obtained models using the following command:

anova(sleepmod1,sleepmod)

This yields the following result:

```
refitting model(s) with ML (instead of REML)
Data: sleepstudy
Models:
sleepmod1: Reaction ~ 1 + Days + (1 | Subject) + (Days - 1 | Subject)
sleepmod: Reaction ~ 1 + Days + (Days | Subject)
      Df    AIC    BIC   logLik deviance Chisq Chi Df Pr(>Chisq)
sleepmod1  5  1762.0 1778.0 -876.00  1752.0
sleepmod   6  1763.9 1783.1 -875.97  1751.9  0.0639  1  0.8004
```

The high p-value of 0.8 indicates that the model assuming dependence is not better than the model assuming independence. This corresponds with the low Chi-square statistic of 0.0639. Thus, we would opt for the smaller model, we assume independence. Truck drivers with high initial reaction times are generally not affected differently by sleep deprivation than truck drivers who reacted quickly at first. This corresponds with the low correlation coefficient of 0.07 in the bigger model assuming dependence.

Please note that R refits both models using ML instead of REML when comparing models using ANOVA.

The anova function also shows the number of parameters in the model (Df), the Akaike Information Criterion (AIC), the Bayesian Information Criterion (BIC), the log-likelihood (logLik) using the estimates and the deviance. The logLik is in this case the value of the full maximum likelihood. The AIC and BIC are also criteria to compare the fitted models. The deviance equals $-2 \cdot \log\text{Lik}$, while the AIC and BIC are given by the following equations:

$$\text{AIC} = \text{deviance} + 2 \cdot \text{Df}$$

$$\text{BIC} = \text{deviance} + \text{Df} \cdot \log(\text{N})$$

When using AIC to compare models for the same data, we prefer the model with the lowest AIC. Similarly, when using BIC we prefer the model with the lowest BIC. Alternatively, the REML criterion may be used to compute an REML version of AIC or BIC. However, this is not the case when using the anova function in R.

4 Muscle building example

The data for the muscle building example was taken from Applied Longitudinal Analysis by Fitzmaurice et al. (2004). The muscle building data is from a study of exercise therapies where 37 patients were assigned to one of two weightlifting programs. In the first program the number of repetitions was increased as subjects became stronger. In the second program the number of repetitions was fixed but the amount of weight was increased as

subjects became stronger. Measures of strength were taken at baseline (day 0), and on days 2, 4, 6, 8, 10, and 12.

In this example we will not be looking at the independence of the random intercept and slope, but instead we will be looking at the necessity of some of the factors. Specifically, we will be looking at the random slope with respect to the variable time and we will be looking at the fixed explanatory variable program. To do this we fit a mixed effects model with linear time dependence in the random effects part. In the fixed effects part of the model we also include interaction with the program. We then compare this model with the 2 variations.

Our first and biggest model will be called using the following command:

```
strmod=lmer(strength~time*program+(time|id),data=musclebuilding)
```

We then obtain the following output:

```
Linear mixed model fit by REML ['lmerMod']  
Formula: strength ~ time * program + (time | id)  
Data: musclebuilding
```

```
REML criterion at convergence: 818.5
```

```
Scaled residuals:
```

```
  Min   1Q   Median   3Q   Max  
-1.9432 -0.6199 -0.0612  0.5352  3.2512
```

```
Random effects:
```

```
Groups Name      Variance Std.Dev. Corr  
id      (Intercept) 9.95309  3.1549  
        time       0.03433  0.1853 -0.03  
Residual          0.66469  0.8153  
Number of obs: 239, groups: id, 37
```

```
Fixed effects:
```

```
          Estimate Std. Error t value  
(Intercept) 80.13241  0.80129  100.00  
time         0.11699  0.05050   2.32  
program2     1.13141  1.06373   1.06  
time:program2 0.05198  0.06745   0.77
```

```
Correlation of Fixed Effects:
```

```
          (Intr) time  prgrm2  
time      -0.084  
program2  -0.753 0.063  
time:prgrm2 0.063 -0.749 -0.085
```

The output shows an intercept and a slope with respect to time for both random and fixed effects. For the fixed effects the output also shows a slope with respect to program and a slope with respect to the interaction between program and time. In addition the output shows the correlations.

Please note that program corresponds only to the subjects that followed program two. Program one, being the base program, is incorporated into the intercept.

For the fixed effects the output shows that the estimate for the intercept is 80.132, the time component has an estimate of 0.117, the program has an estimate of 1.131 and the interaction between time and program has an estimate of 0.052. This would indicate that the base strength (at time 0 and program one) is about 80 and is increased by about 0.12 for each time. Following program two results in an additional 1.13, while the interaction with time increases strength by another 0.05 for each time.

For the random effects the output shows that the intercept for each individual may vary with a standard deviation of about 3.2, while the variation in the slope has a standard deviation of about 0.2. The output also shows a very low dependence between intercept and time with a correlation coefficient of -0.03. While we will not fit a model without dependence in this example, the low correlation coefficient indicates that such a model would probably be appropriate.

Next we call a smaller model that does not include the random slope effect, but the intercept only. We use the following R command:

```
strmod1=lmer(strength~time*program+(1|id),data=musclebuilding)
```

This results in the following:

```
Linear mixed model fit by REML ['lmerMod']  
Formula: strength ~ time * program + (1 | id)  
Data: musclebuilding
```

REML criterion at convergence: 881.2

Scaled residuals:

Min	1Q	Median	3Q	Max
-2.5979	-0.5463	0.0335	0.6079	3.0065

Random effects:

Groups	Name	Variance	Std.Dev.
id	(Intercept)	10.64	3.263
	Residual	1.23	1.109

Number of obs: 239, groups: id, 37

Fixed effects:

	Estimate	Std. Error	t value
(Intercept)	80.11219	0.83792	95.61
time	0.12144	0.02716	4.47
program2	1.20998	1.11229	1.09
time:program2	0.03056	0.03684	0.83

Correlation of Fixed Effects:

(Intr)	time	prgrm2
--------	------	--------


```
time          -0.189
program2     -0.753 0.142
time:prgrm2  0.139 -0.737 -0.188
```

The output shows the same parameters for the fixed effects. However, the estimates are slightly different. For the random effects the output shows an intercept only and thus, no correlation.

For the fixed effects the output shows that the estimate for the intercept is 80.112, the time component has an estimate of 0.121, the program has an estimate of 1.21 and the interaction between time and program has an estimate of 0.031. This would indicate that the base strength is about 80 and is increased by about 0.12 for each time. Following program two results in an additional 1.21, while the interaction with time increases strength by another 0.03 for each time. For the random effects the output shows that the intercept for each individual may vary with a standard deviation of about 3.3, being slightly higher than the bigger model.

The next step is to compare both models to see if the random time slope effect is significant. Here we use the anova command in R as follows:

anova(strmod,strmod1)

```
refitting model(s) with ML (instead of REML)
Data: musclebuilding
Models:
strmod1: strength ~ time * program + (1 | id)
strmod:  strength ~ time * program + (time | id)
      Df  AIC   BIC   logLik deviance Chisq Chi Df Pr(>Chisq)
strmod1 6  884.85 905.71 -436.42  872.85
strmod   8  828.42 856.23 -406.21  812.42   60.428   2 7.553e-14 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

The ANOVA gives a large Chi-square statistic of 60.428 and a really small p-value nearing 0, this shows that the slope effect is significant and we conclude that the random effects on the slopes are necessary. Now we like to know whether the fixed explanatory variable, program, is necessary. Thus we set up a model without the program effect as fixed effect using the following command:

strmod2=lmer(strength~time+(time | id),data=musclebuilding)

The results are the following:

```
Linear mixed model fit by REML ['lmerMod']
Formula: strength ~ time + (time | id)
Data: musclebuilding
```

REML criterion at convergence: 818.8

Scaled residuals:

Min	1Q	Median	3Q	Max
-1.9259	-0.6214	-0.0544	0.5285	3.2314

Random effects:

Groups	Name	Variance	Std.Dev.	Corr
id	(Intercept)	9.98990	3.1607	
	time	0.03371	0.1836	0.00
Residual		0.66520	0.8156	

Number of obs: 239, groups: id, 37

Fixed effects:

	Estimate	Std. Error	t value
(Intercept)	80.77546	0.52795	153.0
time	0.14619	0.03322	4.4

Correlation of Fixed Effects:

	(Intr)
time	-0.061

The output shows an intercept and a slope with respect to time for both random and fixed effects. In addition the output shows the correlation between intercept and slope with respect to time.

For the fixed effects the output shows that the estimate for the intercept is now 80.77 and the time component has an estimate of 0.146. This would indicate that the base strength is about 81 and is increased by about 0.15 for each time. Both estimates are slightly higher compared to the biggest model.

For the random effects the output shows that the intercept for each individual may vary with a standard deviation of about 3.2, while the variation in the slope has a standard deviation of about 0.2. The output also shows a very low dependence between intercept and time with a correlation of zero, indicating again that a model without dependence would probably be appropriate.

Again we use the anova command to compare this model with the biggest model, strmod.

```
> anova(strmod2,strmod)
refitting model(s) with ML (instead of REML)
Data: musclebuilding
Models:
strmod2: strength ~ time + (time | id)
strmod: strength ~ time * program + (time | id)
      Df  AIC   BIC  logLik deviance Chisq Chi Df Pr(>Chisq)
strmod2 6  826.35 847.21 -407.17  814.35
strmod  8  828.42 856.23 -406.21  812.42  1.9304  2  0.3809
```

The anova function gives a p-value of 0.38 and a Chi-squared statistic of 1.93 and this would indicate using the model without program. So after comparing 3 different models we opt for the strmod2 model, which has the following formula:

$$\text{strmod2: strength} \sim \text{time} + (\text{time} \mid \text{id})$$

Thus, the strength of each subject depends on an intercept and a time effect slope for both the fixed as well as the random effects parts.

Discussion

The goal of this paper was to explore both the theory and the application of mixed effects models. The (linear) Mixed Effects model is an expansion of the most basic statistical models, the linear regression model. Often a linear regression model is not capable of capturing all the information provided by data if experiments are getting more complex. Mixed effects models may be useful when analyzing data of certain experiments, depending on the way these experiments are designed. Typical data designs for which mixed effects models prove useful are blocked design, nested design, repeated measures design, and longitudinal design.

As with any parametric modeling methods, estimates of these parameters need to be calculated. A common and well-known estimate is the maximum likelihood estimate. For mixed effects models the most common estimates are the maximum likelihood (ML) estimate and the restricted maximum likelihood (REML) estimate. However, alternatives have been explored, for instance by Li and Pourahmadi (2013).

Statistical programs such as R contain many build-in realizations of statistical procedures concerning linear regression models. This also includes mixed effects models. Using R, the lmer function can be applied on the data. One can model the data and perform several tests, for instance likelihood-ratio tests, to assess the models and obtain the most accurate model for the experiment. We demonstrate the application of this methodology by considering two real data sets from Belenky et al. (2003) and Fitzmaurice et al. (2004).

There exists a large pool of information on mixed effects models that is beyond the scope of this paper, and is rather technical. However, only a limited amount of papers exists on improving current methods regarding mixed effects models. We mentioned one of such papers dedicated to an alternative REML method for mixed effects models.

References

- Bates, D.M. (2010). lme4: Mixed-effects modeling with R. [<http://lme4.r-forge.r-project.org/IMMwR/lrgprt.pdf>]
- Belenky, G., Wesensten, N.J., Thorne, D.R., Thomas, M.L., Sing, H.C., Redmond, D.P., Russo, M.B. and Balkin, T.J. (2003). Patterns of performance degradation and restoration during sleep restriction and subsequent recovery: a sleep dose-response study. *Journal of Sleep Research* 12: 1-12 [<http://spokane.wsu.edu/researchoutreach/sleep/documents/2003JSR-Belenky-et-al.pdf>]
- Fitzmaurice, G.M., Laird, N.M. and Ware, J.H. (2004). *Applied Longitudinal Analysis*. New York: John Wiley and Sons.
- Fox, J. (2000). *Linear Mixed Models: Appendix to An R and S-PLUS Companion to Applied Regression*. [<http://cran.r-project.org/doc/contrib/Fox-Companion/appendix-mixed-models.pdf>]
- Li, E. and Pourahmadi, M. (2013). An alternative REML estimation of covariance matrices in linear mixed models. *Statistics and Probability Letters*, 83, 1071—1077.
- Pinheiro, J.C. and Bates, D.M. (2000). *Mixed effects models in s and S-plus*. New York: Springer
- West, B.T., Welch, K.B., & Galecki, A. T. (2007). *Linear mixed models: A practical guide to using statistical software*. New York: Chapman & Hall/CRC.
- Winter, B. (2013). *Linear models and linear mixed effects models in R with linguistic applications*. arXiv:1308.5499. [<http://arxiv.org/pdf/1308.5499.pdf>]