

# R-Adelaide 2016

## Day 3: Mixed effects models in R

Dr Steven Delean - Benham Rm 103  
[steven.delean@adelaide.edu.au](mailto:steven.delean@adelaide.edu.au)

School of Biological Sciences - The University of Adelaide



COMMONWEALTH OF AUSTRALIA

Copyright Regulations 1969

**WARNING**

This material has been reproduced and communicated to you by or on behalf of The University of Adelaide under Part VII of the Copyright Act 1968 (the Act).

The material in this communication may be subject to copyright under the Act. Any further reproduction or communication of this material by you may be the subject of copyright protection under the Act.

Do not remove this notice.

# Linear mixed effects models


- Linear mixed-effects models are used when you have random effects
  - Occurs when your subjects or units are grouped
  - 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 the random groups)
  - when you have measurements on related individuals (family is the random group)
  - when you measure subjects or units repeatedly (subject or unit is the random group)

# What are we actually talking about?

Alternative names for mixed effects models


- Mixed model, LMM/GLMM
- Random effects model, mixed effects model
- Longitudinal regression, repeated measures model
- Hierarchical model, multilevel model
- Covariance pattern model

# Example of clustered data




# Example of clustered data

- Record students' math achievement scores where students were randomly selected from schools in a state



# Example of longitudinal data

- 18 study participants are sleep deprived for several days
- Each day, measure that subject's reaction time



# Definition of LME Models

In matrix notation a linear model can be represented as

$$\hat{y}_i = \beta_0 + \beta_1 x_1 + \cdots + \epsilon_i$$

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

# Definition of LME Models

In matrix notation an LME model can be represented as

$$y = X\beta + Zb + \epsilon$$

where

- $y$  is a vector of observations, with mean  $E(y) = X\beta$
- $\beta$  is a vector of fixed effects
- $b$  is a vector of random effects with mean  $E(b) = 0$  and variance-covariance matrix  $\text{var}(b) = G$
- $\epsilon$  is a vector of IID random error terms with mean  $E(\epsilon) = 0$  and variance  $\text{var}(\epsilon) = R$
- $X$  and  $Z$  are matrices of regressors relating the observations  $y$  to  $\beta$  and  $b$ , respectively.
- Fixed effects have coefficients
- Random effects have estimates of variation



# Mixed models in R

- Maximum likelihood or restricted maximum likelihood (REML) estimates of the parameters in linear mixed-effects
- Historically, `nlme` was main package for mixed models in R
  - Includes linear and nonlinear fixed effects
  - Variety of variance functions and correlation structures
- Development then shifted to package `lme4`
- The term "mixed effects" denotes model with both fixed- and random-effects terms in a linear predictor

# Mixed model formulas - lmer

- As for linear models, first two arguments to `lmer`
  - Formula specifying the model
  - Data
- Formula for `lm`:

```
1 response ~ expr
```

- Formula for `lmer`:

```
1 response ~ FEexpr + (REexpr1 | factor1) + (REexpr2 | factor2) + ...
```

where,

- `FEexpr` is fixed-effects model matrix,  $X$
- Random-effects terms `(REexpr1 | factor1)` and `(REexpr2 | factor2)` determine random-effects model matrix,  $Z$

# Mixed model formulas - lmer


Formula	Alternative	Meaning
$(1 \mid g)$	$1 + (1 \mid g)$	Random intercept with fixed mean
$0 + \text{offset}(o) + (1 \mid g)$	$-1 + \text{offset}(o) + (1 \mid g)$	Random intercept with a priori means
$(1 \mid g1/g2)$	$(1 \mid g1) + (1 \mid g1:g2)$	Intercept varying among g1 and g2 within g1
$(1 \mid g1) + (1 \mid g2)$	$1 + (1 \mid g1) + (1 \mid g2)$	Intercept varying among g1 and g2
$x + (x \mid g)$	$1 + x + (1 + x \mid g)$	Correlated random intercept and slope
$x + (x \parallel g)$	$1 + x + (1 \mid g) + (0 + x \mid g)$	Uncorrelated random intercept and slope

# A Greenhouse Experiment testing C:N Ratios

- Does changing the C:N Ratio of soil affects plant leaf growth?
  - 3 treatments: control, C addition, and N addition
  - 4 pots per treatment
  - 3 leaves per plant

```
1 plants <- read.csv("data-raw/nestedGrowth.csv")
2 head(plants)
3 # Set Control group as reference level
4 plants$Treatment <- with(plants, relevel(Treatment, ref = "Control"))
5 # Plot data
6 plantPlot <- ggplot(data=plants, aes(x=Pot, y=Growth, color=Treatment,
7   ↪ group=Pot)) + theme_bw()
8 plantPlot + geom_point(size=3)
```

	X	Treatment	Pot	Growth
1	1	Add C	a	6.13
2	2	Add C	a	7.42
3	3	Add C	a	8.20
4	4	Add C	b	15.73
5	5	Add C	b	9.65
6	6	Add C	b	13.14




# Simplifying analysis - averaging

- If design is balanced (and no interest in within pot variance) average within each pot

```
1 stat_sum_single <- function(fun, geom="point", ...) {
2   stat_summary(fun.y=fun, shape=17, geom=geom, size=6, ...)
3 }
4 plantPlot+ geom_point(alpha=0.3) + stat_sum_single(mean)
5 # ANOVA
6 library(car)
7 library(dplyr)
8 meanPlants <- plants %>% group_by(Pot, Treatment) %>% summarise(Growth =
9   ↪ mean(Growth))
10 plantGrowth <- lm(Growth ~ Treatment, data = meanPlants)
   Anova(plantGrowth)
```

```
1 Anova Table (Type II tests)
2
3 Response: Growth
4      Sum Sq Df F value    Pr(>F)
5 Treatment 217.67771  2  8.91584 0.0073312 **
6 Residuals 109.86628  9
7 ---
8 Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```



# Classical ANOVA Error Decomposition with Expected Mean Squares

$$\bullet SS_{Total} = SS_{Treatment} + SS_{PotError} + SS_{WithinPotError}$$

```
1 plantAOV <- aov(Growth ~ Treatment + Error(Pot), data = plants)
2 summary(plantAOV)
```

```
1 Error: Pot
2      Df  Sum Sq Mean Sq F value    Pr(>F)
3 Treatment  2 653.0331  326.5166  8.91584 0.0073312 **
4 Residuals  9 329.5988   36.6221
5 ---
6 Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
7
8 Error: Within
9      Df  Sum Sq Mean Sq F value    Pr(>F)
10 Residuals 24 97.68967  4.070403
11
```

# Multilevel/Clustered/Hierarchical/Mixed Model


- $i$  = treatment,  $j$  = pot,  $k$  = subsample

$$\hat{y}_{ijk} = b_{j[i]} + \beta_i X + \epsilon_{ijk}$$

$$b_{j[i]} \sim \mathcal{N}(\mu_b, \sigma_b^2)$$

$$\epsilon_{ijk} \sim \mathcal{N}(0, \sigma^2)$$

- What is a random effect?
  - Each pot has different parameter value
  - Unlike in linear model, intercept values are constrained around Normal distribution



# Some Points about Multilevel Models

- Flexible. Can accommodate varying slope, intercept, intercept-slope models
- Linear mixed models are solved using Restricted Maximum Likelihood (REML)
  - ML estimation produces downward biased estimates of random effect variances
- As group level effects are drawn from the same distribution, Best Linear Unbiased Predictors (BLUPs) are shrunk towards grand mean - basically, we use information from all groups to inform within group means - useful for unbalanced designs



# Many R Packages for Multilevel Models

- nlme - from Pinhero and Bates 2009
- lmer - bleeding edge by Doug Bates
- MCMCglmm - uses Bayesian techniques & MCMC (similar syntax to nlme)
- glmmADMB - interface for AD Model Builder
- etc

# Fitting a Varying Intercept Model for the Greenhouse Experiment - lmer

```
1 plantLMM <- lmer(Growth ~ Treatment + (1 | Pot), data=plants)
2 summary(plantLMM)
```

```
1 Linear mixed model fit by REML
2 t-tests use Satterthwaite approximations to degrees of freedom ['lmerMod
  ↪ ']
3 Formula: Growth ~ Treatment + (1 | Pot)
4 Data: plants
5
6 REML criterion at convergence: 167.2
7
8 Scaled residuals:
9      Min       1Q   Median       3Q      Max
10 -1.60894241 -0.73479195 -0.08365836  0.76529269  1.63672837
11
12 Random effects:
13  Groups Name      Variance Std.Dev.
14  Pot      (Intercept) 10.850564 3.294019
15  Residual              4.070403 2.017524
16 Number of obs: 36, groups: Pot, 12
17
18 Fixed effects:
19      Estimate Std. Error      df t value    Pr(>|t|)
20 (Intercept)  14.403830   1.746952  8.999998  8.24512 1.7379e-05 ***
21 TreatmentAdd C -5.301297   2.470563  8.999998 -2.14579  0.060450 .
22 TreatmentAdd N  5.130825   2.470563  8.999998  2.07678  0.067617 .
23 ---
24 Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
25
26 Correlation of Fixed Effects:
27      (Intr) TrtmAC
28 TrtmntAddC -0.707
29 TrtmntAddN -0.707  0.500
```

$$\hat{y}_{ijk} = b_{j[i]} + \beta_i X + \epsilon_{ijk}$$

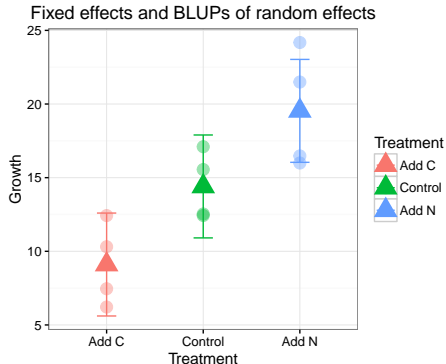
```
1 plantLMM.Ests <- as.data.frame(cbind(round(coef(plantLMM)$Pot, 3), round(
  ↪ ranef(plantLMM)$Pot, 3)))
2 names(plantLMM.Ests)[4] <- "BLUP.Random.Intercepts"
3 print(plantLMM.Ests)
```

	(Intercept)	TreatmentAdd C	TreatmentAdd N	BLUP.Random.Intercepts
1 a	12.759	-5.301	5.131	-1.645
2 b	17.727	-5.301	5.131	3.323
3 c	15.612	-5.301	5.131	1.208
4 d	11.518	-5.301	5.131	-2.886
5 e	12.534	-5.301	5.131	-1.870
6 f	15.552	-5.301	5.131	1.149
7 g	17.100	-5.301	5.131	2.696
8 h	12.429	-5.301	5.131	-1.975
9 i	19.046	-5.301	5.131	4.642
10 j	11.344	-5.301	5.131	-3.060
11 k	16.362	-5.301	5.131	1.958
12 l	10.864	-5.301	5.131	-3.540

# Visualizing Fixed Effects and BLUPs

Note use of altered model for ease of plotting

```
1 plantLMM2 <- lmer(Growth ~ 0 + Treatment + (1 | Pot), data=plants)
2 summary(plantLMM2)
3 # Fixed effect estimates
4 fixDF <- data.frame(summary(plantLMM2)$coef)
5 fixDF$Treatment <- factor(gsub("Treatment", "", rownames(fixDF)), levels
6   ↪ = c("Add C", "Control", "Add N"))
7 # Random effect BLUPs
8 ranDF <- data.frame(Estimate = unlist(ranef(plantLMM, condVar = TRUE)$Pot
9   ↪ ) + rep(fixDF$Estimate[c(2, 1, 3)], rep(4, 3)))
10 ranDF$Treatment <- factor(as.character(meanPlants$Treatment), levels = c(
11   ↪ "Add C", "Control", "Add N"))
12 # Plot
13 p <- ggplot(fixDF, aes(x=Treatment, y=Estimate, color=Treatment)) +
14   theme_bw() +
15   ylab("Growth") +
16   geom_errorbar(aes(ymin=Estimate-Std..Error*2, ymax=Estimate+Std..Error*
17     ↪ 2), width = 0.2) +
18   geom_point(shape = 17, size = 6) +
19   ggtitle("Fixed effects and BLUPs of random effects")
20 p + geom_point(mapping = aes(x=Treatment, y=Estimate), data = ranDF, size
21   ↪ =4, alpha = 0.4)
```





# What next?

- Diagnostics in Mixed Models
- Assessing hypotheses
- Hierarchical Models
- Variable Slope-Intercept models
- A Procedure for Assessing Mixed Model Structure
- A Taste of Generalized Linear Mixed Models



# Residuals diagnostics for observational units

```
1 # Plot residuals vs fitted values
2 plot(fitted(plantLMM), residuals(plantLMM), xlab="Fitted values", ylab="Residuals", pch=21, bg="lightblue")
3 abline(a=0, b=0, col="black", lty=2)
4 # Q-Q plot of residuals
5 qqnorm(residuals(plantLMM), type="pearson", main="", pch=21, bg="lightblue")
6 qqline(residuals(plantLMM), type="pearson")
```



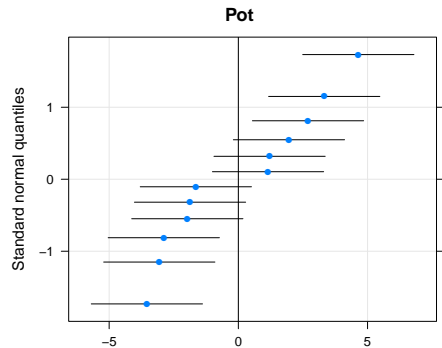
# Residuals diagnostics at group-level

```
1 # Plot residuals vs fitted values at group level
2 plot(predict(plantLMM2, re.form=-0), plants$Growth - predict(plantLMM2, re.form=-0), xlab="Group-level fitted values", ylab="Group-level residuals"
   ↪ , pch=21, bg="lightblue")
3 abline(a=0, b=0, col="black", lty=2)
4 # Q-Q plot of random effects
5 r <- unlist(ranef(plantLMM, condVar = TRUE)$Pot)
6 qqnorm(r/sd(r), main="", pch=21, bg="lightblue")
7 qqline(r/sd(r))
```



# Random effects BLUPS

```
1 library(lattice)
2 qqmath(ranef(plantLMM, condVar = TRUE), strip = FALSE)
3 # Test for heterogeneity
4 library(RLRsim)
5 exactRLRT(plantLMM)
```



# Model inference

```
1 # Global testing for fixed effects
2 drop1(plantLMM, test = "Chisq")
3 # Using contrasts for planned comparisons
4 library(lsmmeans)
5 trt.lsm <- lsmeans(plantLMM, "Treatment")
6 contrast(trt.lsm, "trt.vs.ctrl", adjust = "none")
7 # Calculating bootstrap confidence intervals - parameter contrasts
8 confint(plantLMM, method = "boot", boot.type = "basic", nsim = 100,
9        ↪ parallel = "multicore", ncpus = 8)
10 # Calculating bootstrap confidence intervals - Treatment means
11 confint(plantLMM2, method = "boot", boot.type = "basic", nsim = 100,
12        ↪ parallel = "multicore", ncpus = 8)
```

## Single term deletions

Model:

Growth ~ Treatment + (1 | Pot)

	Df	AIC	LRT	Pr(Chi)
<none>		185.60902		
Treatment	2	194.71732	13.1083	0.0014242 **

---

Signif. codes: 0 '\*\*\*' 0.001 '\*\*' 0.01 '\*' 0.05 '.' 0.1 ' ' 1


		2.5 %	97.5 %
.sig01		1.06733677269	4.920323813
.sigma		1.34597084395	2.635024473
(Intercept)		11.74738290407	17.594991016
TreatmentAdd C	-10.70452432258	-1.196672772	
TreatmentAdd N	0.01244979266	9.969124979	

		2.5 %	97.5 %
.sig01		1.359105917	4.849531572
.sigma		1.411892016	2.549262368
TreatmentControl	10.245785828	18.000970897	
TreatmentAdd C	6.097065684	12.975063104	
TreatmentAdd N	15.193178109	24.386497203	



# Example dataset: sleepstudy


- 18 study participants are sleep deprived for several days
- Each day, measure that subject's reaction time



# A simple linear model

```
1 #*BEGIN_SRC R :results output raw :exports both
2 lm.simple <- lm(Reaction ~ Days, data=sleepstudy)
3 summary(lm.simple)
```


```
1 Call:
2 lm(formula = Reaction ~ Days, data = sleepstudy)
3
4 Residuals:
5     Min       1Q   Median       3Q      Max
6  -110.85  -27.48   1.55   26.14  139.95
7
8 Coefficients:
9             Estimate Std. Error t value Pr(>|t|)
10 (Intercept)    251.41         6.61  38.03  < 2e-16 ***
11 Days             10.47         1.24   8.45  9.9e-15 ***
12 ---
13 Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
14
15 Residual standard error: 47.7 on 178 degrees of freedom
16 Multiple R-squared:  0.286,    Adjusted R-squared:  0.282
17 F-statistic: 71.5 on 1 and 178 DF,  p-value: 9.89e-15
```



# Fitting a varying intercepts model with `lm`

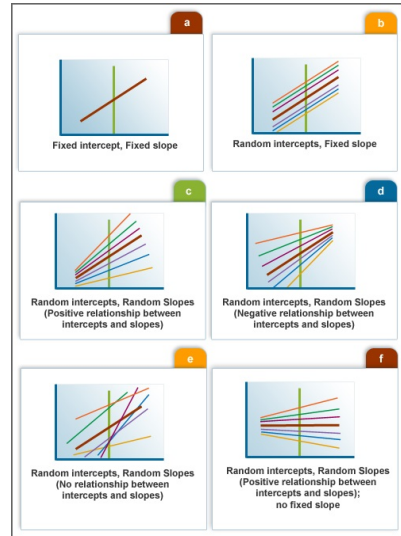
```
1 #*BEGIN_SRC R :results output raw :exports both
2 lm.complete <- lm(Reaction ~ 0 + factor(Subject) + Days, data=sleepstudy)
3 summary(lm.complete)
```

```
1
2 Call:
3 lm(formula = Reaction ~ 0 + factor(Subject) + Days, data = sleepstudy)
4
5 Residuals:
6     Min       1Q   Median       3Q      Max
7  -100.54  -16.39   -0.34   15.22   131.16
8
9 Coefficients:
10              Estimate Std. Error t value Pr(>|t|)
11 factor(Subject)308    295.031    10.447   28.2   <2e-16 ***
12 factor(Subject)309    168.130    10.447   16.1   <2e-16 ***
13 factor(Subject)310    183.898    10.447   17.6   <2e-16 ***
14 factor(Subject)330    256.119    10.447   24.5   <2e-16 ***
15 factor(Subject)331    262.333    10.447   25.1   <2e-16 ***
16 factor(Subject)332    260.199    10.447   24.9   <2e-16 ***
17 factor(Subject)333    269.056    10.447   25.8   <2e-16 ***
18 factor(Subject)334    248.199    10.447   23.8   <2e-16 ***
19 factor(Subject)335    202.967    10.447   19.4   <2e-16 ***
20 factor(Subject)337    328.618    10.447   31.5   <2e-16 ***
21 factor(Subject)349    228.732    10.447   21.9   <2e-16 ***
22 factor(Subject)350    266.500    10.447   25.5   <2e-16 ***
23 factor(Subject)351    242.995    10.447   23.3   <2e-16 ***
24 factor(Subject)352    290.319    10.447   27.8   <2e-16 ***
25 factor(Subject)369    258.932    10.447   24.8   <2e-16 ***
26 factor(Subject)370    244.599    10.447   23.4   <2e-16 ***
27 factor(Subject)371    247.881    10.447   23.7   <2e-16 ***
28 factor(Subject)372    270.783    10.447   25.9   <2e-16 ***
29 Days
30      10.467      0.804   13.0   <2e-16 ***
---
```



# Mixed models enable us to account for variability

- Varying intercepts
- Varying slopes




## Hence there's gradient between

- Mixed models estimate varying parameters (intercepts and/or slopes) with pooling among levels (rather than considering them fully independent)
- **complete pooling**: Single overall intercept
  - `lm (Reaction ~ Days)`
- **no pooling**: One **independent** intercept for each plot
  - `lm (Reaction ~ Days + factor(Subject))`
- **partial pooling**: Inter-related intercepts
  - `lm (Reaction ~ Days + (1 | Subject))`

# Fitting a varying intercepts model with lmer

```
1 #*BEGIN_SRC R :results output raw :exports both
2 fm.rInt <- lmer(Reaction ~ Days + (1 | Subject), sleepstudy)
3 summary(fm.rInt)
```


```
1 Linear mixed model fit by REML
2 t-tests use Satterthwaite approximations to degrees of freedom ['lmerMod
  ↪ ']
3 Formula: Reaction ~ Days + (1 | Subject)
4 Data: sleepstudy
5
6 REML criterion at convergence: 1786.5
7
8 Scaled residuals:
9      Min       1Q   Median       3Q      Max
10 -3.226 -0.553  0.011  0.519  4.251
11
12 Random effects:
13  Groups   Name      Variance Std.Dev.
14  Subject (Intercept) 1378     37.1
15  Residual                960     31.0
16 Number of obs: 180, groups: Subject, 18
17
18 Fixed effects:
19             Estimate Std. Error    df t value Pr(>|t|)
20 (Intercept)  251.405     9.747 22.800   25.8   <2e-16 ***
21 Days         10.467     0.804 161.000   13.0   <2e-16 ***
22 ---
23 Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
24
25 Correlation of Fixed Effects:
26      (Intr)
27 Days -0.371
```



# Fitting a varying intercepts and slopes model with lmer

```
1 fm.rIntSlp <- lmer(Reaction ~ Days + (Days | Subject), sleepstudy)
2 summary(fm.rIntSlp)
```

```
1 Linear mixed model fit by REML
2 t-tests use Satterthwaite approximations to degrees of freedom ['lmerMod
  ↪ '']
3 Formula: Reaction ~ Days + (Days | Subject)
4 Data: sleepstudy
5
6 REML criterion at convergence: 1743.6
7
8 Scaled residuals:
9     Min       1Q   Median       3Q      Max
10  -3.954  -0.463   0.023   0.463   5.179
11
12 Random effects:
13  Groups   Name      Variance Std.Dev. Corr
14  Subject (Intercept) 612.1    24.74
15           Days       35.1     5.92   0.07
16  Residual          654.9    25.59
17 Number of obs: 180, groups: Subject, 18
18
19 Fixed effects:
20             Estimate Std. Error    df t value Pr(>|t|)
21 (Intercept)  251.41      6.82  17.00  36.84 < 2e-16 ***
22 Days         10.47      1.55  17.00   6.77 3.3e-06 ***
23 ---
24 Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
25
26 Correlation of Fixed Effects:
27      (Intr)
28 Days -0.138
```



# Why no p-values in the `lme4` summaries?

- Recall that p-values in the coefficient summary of fitted linear models are the probability of getting a test statistic as large (or larger) if the coefficient was indeed 0.
- Also recall that p-values are determined using null reference distributions
  - Under the null hypothesis, the test statistic has a known distribution
- In LME models, the null reference distribution is technically not known, at least not for unbalanced data
  - Thus the `lme4` authors elected to not output p-values based on a distribution that is not actually the distribution of the test statistic



# How do I know if a coefficient is "significant"?

- Quick, common-sense way: look at the  $t$  value
- If the  $t$  value is greater than 2, then the coefficient is likely significant
  - In other words, the coefficient estimate is more than 2 standard errors away from 0
- Note:
  - $t \text{ value} = \text{Estimate} / \text{Std. Error}$
  - Of course ask yourself if the result is practically significant as well

# How do I know if a coefficient is "significant"?

- Another way is to compute 95% confidence intervals using the `confint()` function
  - If the interval contains 0, it is not significant
  - Further, confidence intervals give an indication of coefficient size and variability
- Say your fitted model object is `lme1` Two ways to compute confidence intervals are as follows:
  - `confint(lme1)` computes a likelihood profile and finds the appropriate cutoffs based on the likelihood ratio test
  - `confint(lme1, method="boot")` parametric bootstrapping ( $B = 500$ ) with confidence intervals computed from the bootstrap distribution

# Confidence intervals for the sleepstudy data

```
1 fml <- lmer(Reaction ~ Days + (Days|Subject), sleepstudy)
2 fml.ci <- confint(fml, method="boot")
3 fml.ci
```

```
1          2.5 % 97.5 %
2 .sig01    11.779 34.835
3 .sig02    -0.508  0.998
4 .sig03     3.320  8.213
5 .sigma    22.349 28.566
6 (Intercept) 238.046 265.029
7 Days       7.416 13.407
```

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

REML criterion at convergence: 1743.6

```
Scaled residuals:
   Min      1Q  Median      3Q      Max
-3.954 -0.463  0.023  0.463  5.179
```

```
Random effects:
Groups   Name             Variance Std.Dev. Corr
Subject (Intercept) 612.1      24.74
          Days         35.1       5.92    0.07
Residual          654.9      25.59
Number of obs: 180, groups: Subject, 18
```

```
Fixed effects:
              Estimate Std. Error t value
(Intercept)   251.41      6.82    36.8
Days           10.47      1.55     6.8
```

```
Correlation of Fixed Effects:
      (Intr)
Days -0.138
```

# But I want p-values!

- The `lmerTest` package provides the kind of p-values SAS provides (based on Satterthwaite's approximations)
- Just load it `library(lmerTest)`, run `lmer` as usual and call `summary` on the object

```
1 library(lmerTest)
2 m <- lmer(dv ~ x1 + (x1 | g), data=df)
3 summary(m)
```

A column of p-values is included in the summary output. Again, they're approximate.

## Exercise: Random Effects on Richness using glmer

- Poisson distribution is widely used for the description of count data
- Invertebrates species richness in 45 sites on Dutch coast beaches versus height above average sea level
- We will:
  - Fit data from RIKZ survey
  - Random Effect of Beach ONLY
  - Visualize Random Effects

# Fitting and summarising a mixed model

```
1 rikz <- read.csv("../data-raw/rikz.csv")
2 rikz$Beach <- as.factor(rikz$Beach)
3 rikz$angleRad <- rikz$angle1*pi/180
4 rikz$angleRad
5 mod1_glmmer <- glmer(Richness ~ NAP*angleRad + (1|Beach), data=rikz,
6   ↪ family="poisson")
7 summary(mod1_glmmer)
8 anova(mod1_glmmer, type = 'm')
```

```
1 Generalized linear mixed model fit by maximum likelihood (Laplace
2   ↪ Approximation) ['glmerMod']
```

```
3 Family: poisson ( log )
4 Formula: Richness ~ NAP * angleRad + (1 | Beach)
5 Data: rikz
```

	AIC	BIC	logLik	deviance	df.resid
	217.2	226.2	-103.6	207.2	40

```
9 Scaled residuals:
```

	Min	1Q	Median	3Q	Max
	-1.747	-0.629	-0.212	0.273	2.624

```
12 Random effects:
```


Groups	Name	Variance	Std.Dev.
Beach	(Intercept)	0.307	0.554

Number of obs: 45, groups: Beach, 9

```
18 Fixed effects:
```


	Estimate	Std. Error	z value	Pr(> z )
(Intercept)	1.850	0.216	8.57	< 2e-16 ***
NAP	-0.767	0.135	-5.70	1.2e-08 ***
angleRad	-0.228	0.105	-2.18	0.030 *
NAP:angleRad	0.184	0.083	2.22	0.026 *

---



# Parametric bootstrap to estimate variation in random effects

```
1 # Random intercepts using sim from the arm package
2 library(arm)
3 n.sim <- 1000
4 simu <- sim(mod1_glm, n.sims=n.sim)
5 # Intercept
6 ranInt <- as.data.frame(t(apply(simu@ranef$Beach[, , 1], 2, quantile,
7   ↪ prob=c(0.025, 0.5, 0.975))))
8 names(ranInt) <- c("lwr", "fit", "upr")
9 ranInt$Beach <- paste("b", 1:9, sep = "")
```



# Parametric bootstrap confidence intervals

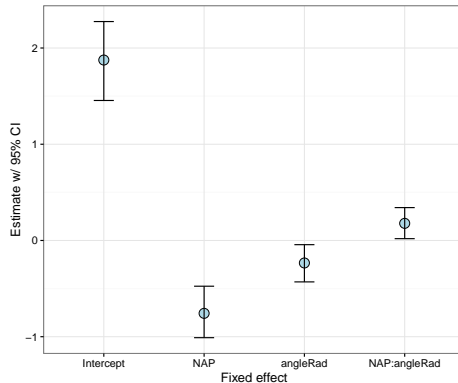
```
1 # mod1_glmer_ci <- confint(mod1_glmer, method="boot")  
2 mod1_glmer_ci
```

```
1          2.5 %   97.5 %  
2 .sig01      0.1635 0.80056  
3 (Intercept) 1.4111 2.27656  
4 NAP        -1.0454 -0.50724  
5 angleRad    -0.5559 -0.00781  
6 NAP:angleRad -0.0415 0.40827
```




# Parametric bootstrap fixed effects estimates

```
1 # Predict fixed effects
2 coeff <- as.data.frame(t(apply(simu@fixef, 2, quantile, prob=c(0.025,
  ↪ 0.5, 0.975))))
3 names(coeff)[1:3] <- c('lwr', 'fit', 'upr')
4 coeff$term <- row.names(coeff)
5 coeff$term[1] <- "Intercept"
6 coeff$term <- factor(coeff$term, levels=c("Intercept", "NAP", "angleRad",
  ↪ "NAP:angleRad"))
7 summary(coeff)
```



# Parametric bootstrap model predictions

```
1 # plotting the effect of NAP on the richness
2 nsim <- 1000
3 bsim <- sim(mod1_glm, n.sim=nsim)
4 newdat <- data.frame(NAP=seq(-1.5, 2.5, length=100), angleRad=mean(rikz$
  ↪ angleRad))
5 Xmat <- model.matrix(~NAP*angleRad, data=newdat)
6 predmat <- matrix(ncol=nsim, nrow=nrow(newdat))
7 predmat <- apply(bsim@fixef, 1, function(x) exp(Xmat*%*%x))
8 newdat$lower <- apply(predmat, 1, quantile, prob=0.025)
9 newdat$upper <- apply(predmat, 1, quantile, prob=0.975)
10 newdat$Richness <- apply(predmat, 1, quantile, prob=0.5)
```



# Differences between lmer and lme

- Package nlme contains function lme
- Package lme4 contains function lmer
- lme has more flexibility in specifying the covariance structure
- lmer allows you to fit GLMMs (using glmer)
  - Includes glmer.nb for negative binomial variance models
- Different specification of random effects in call

```
1 # lmer
2 fm.rIntSlp <- lmer(Reaction ~ Days + (Days | Subject), sleepstudy)
3 summary(fm.rIntSlp)
4 # lme
5 fm.rIntSlp <- lme(Reaction ~ Days, random = ~ Days | Subject, sleepstudy)
6 summary(fm.rIntSlp)
```

# Challenge

- Growth and maternal effects in mice (`mouse.csv`)
  - In a laboratory experiment, 20 female mice each gave birth to a litter of pups
  - Each litter was reared in a single cage, and each pup was weighed at several timepoints
- Some biological questions of interest are:
  - Do males and females differ with respect to their average growth rates?
  - Is there a maternal effect? (i.e., do individuals from the same litter tend to be similar?)
- See R script for data and code!