

# Statistics for EES

## Linear mixed-effects models

Dirk Metzler

July 8, 2020

### Contents

1 Individual differences in exploratory great tits	1
2 The random effect of the beach	6
3 How to select a mixed-effects model	12

## 1 Individual differences in exploratory great tits

In slides:  
image taken from <https://commons.wikimedia.org/wiki/File:Tihane.jpg> (creative commons license)

Great Tit  
*Parus major*  
(German: Kohlmeise)

### References

[1] Niels J. Dingemanse, M. Moiron, Y.G. Araya-Ajoy, A. Mouchet, R.N. Abbey-Lee (2020) Individual variation in age-dependent reproduction: Fast explorers live fast but senesce young? *J Anim Ecol.* 89:601–613

- 12 plots of 50 nest boxes observed for 7 years
- 1209 observations with 813 birds, including e.g.:
  - clutch size and number of fledglings in a year
  - how exploratory is a bird (experiment)
- Main question: do more exploratory birds have their reproductive phase earlier in life?
- Here we focus on age-dependence of exploratory behavior

```
> str(dat)
'data.frame': 1209 obs. of 16 variables:
 $ BirdID      : Factor w/ 813 levels "256","271","273",...: 1 1 1 1 2 3 4 4 4 5 ...
 $ Plot        : Factor w/ 12 levels "10","11","12",...: 11 11 11 11 1 1 7 7 7 7 ...
 $ BroodID     : int  98 533 1702 2902 129 319 207 438 1507 210 ...
 $ Year        : Factor w/ 7 levels "2010","2011",...: 1 2 3 4 1 2 1 2 3 1 ...
 $ MassSummer  : num  18.1 18.1 17.9 18.4 18.3 17.2 17.5 16.5 16.5 18.1 ...
 $ ExpFieldScore: int  43 52 50 58 71 67 62 58 61 57 ...
 $ ClutchSize  : int  7 9 8 11 9 11 7 7 8 11 ...
```

```

$ NumberFledged: int  0 6 5 0 0 10 4 7 7 9 ...
$ AvgDay14Mass  : num  NA 17.4 17.3 14.1 12.8 ...
$ PredTreat     : int  0 0 0 2 0 0 0 0 0 0 ...
$ AgeBreeder    : int  0 1 2 3 0 1 0 1 2 0 ...
$ MinAge       : int  0 0 0 0 0 1 0 0 0 0 ...
$ MaxAge       : int  3 3 3 3 0 1 2 2 2 0 ...
$ Local        : int  0 0 0 0 0 0 0 0 0 0 ...
$ Dead         : int  1 1 1 1 1 1 1 1 1 1 ...
$ BSManip      : int  0 3 0 0 1 1 0 0 0 1 ...

```

For observation  $i$  (line  $i$  in data frame) we denote as

$b_i$  the bird ID (factor BirdID with 813 levels)

$a_i$  age of the bird at that time (numerical AgeBreeder)

$X_i$  how explorative was the bird at that time (numerical ExpFieldScore)

$y_i$  year of observation (factor Year with 7 levels)

$p_i$  plot to which the bird's nest belonged (factor Plot with 12 levels)

First model (only fixed effects):

$$X_i = \delta + \alpha \cdot a_i + \beta_{b_i} + \gamma_{y_i} + \rho_{p_i} + \varepsilon_i$$

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

needs to estimate  $1 + 1 + 812 + 6 + 11 + 1 = 832$  parameters

```

> mod.fix <- lm(ExpFieldScore ~ AgeBreeder + BirdID + Plot + Year, dat)
> drop1(mod.fix, test="F")
Single term deletions

```

Model:

```

ExpFieldScore ~ AgeBreeder + BirdID + Plot + Year
      Df Sum of Sq  RSS   AIC F value    Pr(>F)
<none>                101537 6762.7
AgeBreeder    0             0 101537 6762.7
BirdID       779   376408 477945 6992.3  1.6941 1.054e-08 ***
Plot         1         222 101759 6763.2  0.7779  0.378394
Year        5         5428 106965 6812.8  3.8062  0.002264 **
---

```

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

```
> summary(mod.fix)
```

```

Call:
lm(formula = ExpFieldScore ~ AgeBreeder + BirdID + Plot + Year,
    data = dat)

```

```

Residuals:
    Min       1Q   Median       3Q      Max
-62.488   0.000   0.000   1.069  47.140

```

Coefficients: (11 not defined because of singularities)

```

      Estimate Std. Error t value Pr(>|t|)
(Intercept)  6.005e+01  8.780e+00   6.840 3.46e-11 ***
AgeBreeder  -4.610e+00  8.937e-01  -5.158 4.15e-07 ***
BirdID271    1.095e+01  1.903e+01   0.575 0.565622
BirdID273    1.815e+01  1.896e+01   0.957 0.339122
BirdID276    6.316e+00  1.292e+01   0.489 0.625154
BirdID277   -3.054e+00  1.903e+01  -0.160 0.872600
BirdID291    1.700e+01  1.467e+01   1.159 0.247149
BirdID314   -4.855e+00  3.064e+01  -0.158 0.874207
BirdID318    6.001e+00  1.467e+01   0.409 0.682660
BirdID321    2.946e+00  1.903e+01   0.155 0.877105
BirdID323   -1.018e+00  1.292e+01  -0.079 0.937252
BirdID324    2.315e+01  1.896e+01   1.221 0.222934

```

```

BirdID344  1.095e+01  1.903e+01  0.575 0.565622
BirdID345 -1.805e+01  1.903e+01 -0.949 0.343503
BirdID349 -2.005e+01  1.903e+01 -1.054 0.292779
BirdID363 -1.405e+01  1.903e+01 -0.738 0.460769
BirdID364  3.316e+00  1.292e+01  0.257 0.797546
BirdID367 -3.405e+01  1.903e+01 -1.789 0.074445
.
.
.
.
.
.
.
.
.
.
BirdID18269 2.556e+00  1.899e+01  0.135 0.892997
BirdID18583 8.946e+00  1.903e+01  0.470 0.638661
Plot11 -2.120e+01  2.404e+01 -0.882 0.378394
Plot12 NA NA NA NA
Plot13 NA NA NA NA
Plot14 NA NA NA NA
Plot15 NA NA NA NA
Plot16 NA NA NA NA
Plot17 NA NA NA NA
Plot18 NA NA NA NA
Plot19 NA NA NA NA
Plot20 NA NA NA NA
Plot21 NA NA NA NA
Year2011 -6.593e+00  3.012e+00 -2.189 0.029235 *
Year2012  2.313e+00  3.002e+00  0.770 0.441682
Year2013 -5.277e+00  3.043e+00 -1.734 0.083735
Year2014  9.874e-01  3.040e+00  0.325 0.745564
Year2015 -7.741e-01  3.095e+00 -0.250 0.802666
Year2016 NA NA NA NA
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 16.89 on 356 degrees of freedom
(55 observations deleted due to missingness)
Multiple R-squared:  0.8028, Adjusted R-squared:  0.3615
F-statistic: 1.819 on 797 and 356 DF,  p-value: 1.133e-10

```

## mixed-effects model

$$X_i = \delta + \alpha \cdot a_i + \beta_{b_i} + \gamma_{y_i} + \rho_{p_i} + \varepsilon_i$$

$\beta_1, \dots, \beta_{813} \sim \mathcal{N}(0, \sigma_\beta^2)$    
 $\gamma_{2010}, \dots, \gamma_{2016} \sim \mathcal{N}(0, \sigma_\gamma^2)$    
 $\rho_1, \dots, \rho_{12} \sim \mathcal{N}(0, \sigma_\rho^2)$    
 $\varepsilon_1, \dots, \varepsilon_{1154} \sim \mathcal{N}(0, \sigma^2)$   
needs to estimate 6 parameters:  $\delta, \alpha, \sigma_\beta, \sigma_\gamma, \sigma_\rho, \sigma$

In R formula notation:

`ExpFieldScore ~ AgeBreeder + (1|BirdID) + (1|Year) + (1|Plot)`

```

> library(lme4)
> mod <- lmer(ExpFieldScore ~ AgeBreeder + (1|BirdID) + (1|Year) + (1|Plot), data=dat, subset= !is.na(ExpFieldScore) & !is.na(AgeBreeder))
> summary(mod)
Linear mixed model fit by REML [EigenMod]
Formula: ExpFieldScore ~ AgeBreeder + (1 | BirdID) + (1 | Year) + (1 | Plot)
Data: dat
Subset: !is.na(ExpFieldScore) & !is.na(AgeBreeder)

REML criterion at convergence: 10204.9

Scaled residuals:
   Min       1Q   Median       3Q      Max
-3.9693 -0.5190 -0.0038  0.5368  2.8090

Random effects:
 Groups Name      Variance Std.Dev.
 BirdID (Intercept) 141.151  11.881
 Plot   (Intercept)  10.732   3.276
 Year   (Intercept)   8.154   2.856
 Residual                277.688  16.664
Number of obs: 1154, groups: BirdID, 791; Plot, 12; Year, 7

Fixed effects:
             Estimate Std. Error t value
(Intercept)  71.0436    1.6378    43.38
AgeBreeder   -3.1450    0.6265    -5.02

Correlation of Fixed Effects:
      (Intr)
AgeBreeder -0.245

```

## Is dependence on age significant?

```

> mod0 <- lmer(ExpFieldScore ~ (1|BirdID) + (1|Year) + (1|Plot),
+             data=dat, subset= !is.na(ExpFieldScore) & !is.na(AgeBreeder))
> library(pbkrtest)

```

```

> KRmodcomp(mod,mod0)
F-test with Kenward-Roger approximation; computing time: 2.26 sec.
large : ExpFieldScore ~ AgeBreeder + (1 | BirdID) + (1 | Year) + (1 | Plot)
small : ExpFieldScore ~ (1 | BirdID) + (1 | Year) + (1 | Plot)
      stat      ndf      ddf F.scaling  p.value
Ftest  24.625    1.000 1014.820        1 8.162e-07 ***
---
Signif. codes:  0 *** 0.001 ** 0.01 * 0.05 . 0.1  1
>

```

Is bird-ID effect significant?

```

> mod.nb <- lmer(ExpFieldScore ~ AgeBreeder + (1|Year) + (1|Plot),
+               data=dat, subset= !is.na(ExpFieldScore) & !is.na(AgeBreeder))
> anova(mod.nb,mod)
refitting model(s) with ML (instead of REML)
Data: dat
Subset: !is.na(ExpFieldScore) & !is.na(AgeBreeder)
Models:
mod.nb: ExpFieldScore ~ AgeBreeder + (1 | Year) + (1 | Plot)
mod:    ExpFieldScore ~ AgeBreeder + (1 | BirdID) + (1 | Year) + (1 | Plot)
      Df  AIC  BIC  logLik deviance  Chisq Chi Df Pr(>Chisq)
mod.nb  5 10281 10307 -5135.7   10271
mod     6 10221 10251 -5104.3   10209 62.824    1 2.26e-15 ***
---
Signif. codes:  0 *** 0.001 ** 0.01 * 0.05 . 0.1  1

```

Is bird-ID effect significant?

```

> PBmodcomp(mod,mod.nb)      ## more robust but takes a while
Parametric bootstrap test; time: 109.31 sec; samples: 1000 extremes: 0;
Requested samples: 1000 Used samples: 570 Extremes: 0
large : ExpFieldScore ~ AgeBreeder + (1 | BirdID) + (1 | Year) + (1 | Plot)
small : ExpFieldScore ~ AgeBreeder + (1 | Year) + (1 | Plot)
      stat df  p.value
LRT    62.824  1 2.22e-15 ***
PBtest 62.824   0.001751 **
---
Signif. codes:  0 *** 0.001 ** 0.01 * 0.05 . 0.1  1


```

### Parametric Bootstrapping

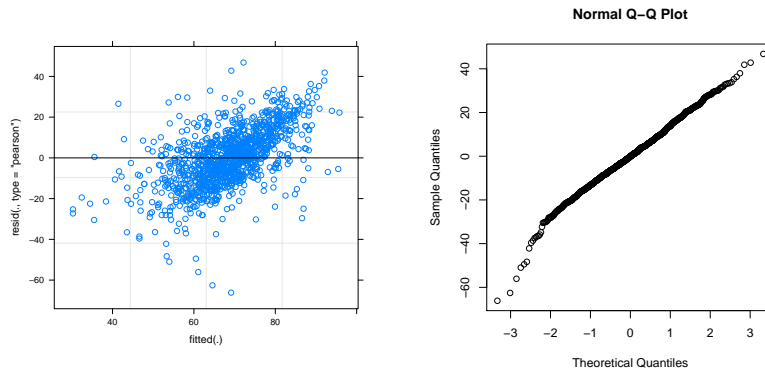
Is model `mod` significantly better than a null model `mod0`?

- Fit `mod` and `mod0` to the data; let  $L$  be the likelihood ratio of `mod` vs. `mod0`.
- repeat  $N$  times (e.g.  $N = 1000$ ):
  - simulate data from `mod0` (with parameters as fitted to original data).
  - fit both models `mod` and `mod0` to the simulated data and calculate the likelihood ratio  $L_i$  for each iteration  $i$ .
- p-value: fraction of  $(L, L_1, L_2, \dots, L_N)$  that are at least as large as  $L$ .

```
library(lattice)
qm <- qqmath(ranef(mod, condVar=TRUE))
qm[[1]]; qm[[2]]; qm[[3]]
```



```
plot(mod)
qqnorm(residuals(mod))
```



Does age-dependence of exploring vary among birds?

$$X_i = \delta + \alpha \cdot a_i + \beta_{b_i} + \phi_{b_i} \cdot a_i + \gamma_{y_i} + \rho_{p_i} + \varepsilon_i$$

$$\phi_1, \dots, \phi_{813} \sim \mathcal{N}(0, \sigma_\phi^2)$$

```
> mod2 <- lmer(ExpFieldScore ~ AgeBreeder + (1 + AgeBreeder || BirdID) + (1|Year) + (1|Plot),
+             data=dat, subset= !is.na(ExpFieldScore) & !is.na(AgeBreeder))
> PBmodcomp(mod2,mod) ## takes a while
Parametric bootstrap test; time: 139.38 sec; samples: 1000 extremes: 2;
Requested samples: 1000 Used samples: 658 Extremes: 2
large : ExpFieldScore ~ AgeBreeder + ((1 | BirdID) + (0 + AgeBreeder | BirdID)) + (1 | Year) + (1 | Plot)
small : ExpFieldScore ~ AgeBreeder + (1 | BirdID) + (1 | Year) + (1 | Plot)
      stat df p.value
LRT    7.9917  1 0.004699 **
PBtest 7.9917  1 0.004552 **
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
Warnings:
1: In checkConv(attr("opt", "derivs"), opt$par, ctrl = control$checkConv, :
  unable to evaluate scaled gradient
2: In checkConv(attr("opt", "derivs"), opt$par, ctrl = control$checkConv, :
  Model failed to converge: degenerate Hessian with 1 negative eigenvalues
```

## 2 The random effect of the beach

Image in slides:

[https://commons.wikimedia.org/wiki/File:Wadden\\_sea\\_and\\_a\\_silhouette\\_of\\_the\\_Maria\\_church\\_in\\_Wierum.jpg](https://commons.wikimedia.org/wiki/File:Wadden_sea_and_a_silhouette_of_the_Maria_church_in_Wierum.jpg)

(creative commons license)

We revisit the RIKZ dataset.

Species abundance and many other variables were measured at 9 beaches.

On every beach, 5 plots were sampled in the intertidal range.

Each plot was sampled only once. Thus, each line in the data table corresponds to one plot.

```
> rikz <- read.csv("RIKZ_Tab4.csv")
> rikz$Beach <- factor(rikz$Beach)
> rikz$Week <- factor(rikz$Week)
> str(rikz)
'data.frame': 45 obs. of 15 variables:
 $ ShannonW      : num  0.76 0.72 0.85 0.53 0.74 0.13 0.4 0.29 1.02 1 ...
 $ Richness      : int   11 10 13 11 10 8 9 8 19 17 ...
 $ angle1       : int   32 62 65 55 23 129 126 52 26 143 ...
 $ angle2       : int   96 96 96 96 96 89 89 89 89 89 ...
 $ exposure     : int   10 10 10 10 10 8 8 8 8 8 ...
 $ salinity     : num  29.4 29.4 29.4 29.4 29.4 29.6 29.6 29.6 29.6 ...
 $ temperature  : num  17.5 17.5 17.5 17.5 17.5 20.8 20.8 20.8 20.8 ...
 $ NAP         : num   0.045 -1.036 -1.336 0.616 -0.684 ...
 $ penetrability: num  254 227 237 249 252 ...
 $ grainsize   : num  222 200 194 221 202 ...
 $ humus       : num   0.05 0.3 0.1 0.15 0.05 0.1 0.1 0.1 0.15 0 ...
 $ chalk       : num   2.05 2.5 3.45 1.6 2.45 2.5 1.85 1.7 2.3 2.6 ...
 $ sorting1    : num  69.8 59 59.2 67.8 57.8 ...
 $ Beach       : Factor w/ 9 levels "1","2","3","4",...: 1 1 1 1 1 2 2 2 2 ...
 $ Week       : Factor w/ 4 levels "1","2","3","4": 1 1 1 1 1 1 1 1 1 ...
```

Today we use the following columns:

**ShannonW** A biodiversity index based on the Shannon entropy.

**Richness** The number of species.

**Beach** Which beach.

**NAP** The altitude of the plot compared to the average sea level.

**exposure** An index composed of wave action, length of the surf zone, slope, grain size, depth of the anaerobic layer.

**salinity** Salt contents.

**humus** Fraction of organic material.

In an earlier analysis, NAP turned out to have an impact on species richness. Does ShannonW also depend on NAP?

```
> mod1 <- lm(ShannonW~NAP,data=rikz)
> summary(mod1)
[. .]
Coefficients:
              Estimate Std. Error t value Pr(>|t|)
(Intercept)  0.47141     0.03627   12.996 < 2e-16 ***
NAP         -0.22583     0.03478   -6.493 7.01e-08 ***
```


---

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

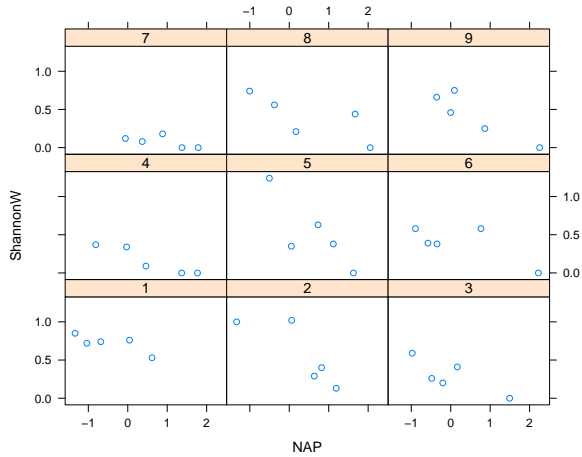
Residual standard error: 0.2294 on 43 degrees of freedom  
Multiple R-squared: 0.4951, Adjusted R-squared: 0.4833  
F-statistic: 42.16 on 1 and 43 DF, p-value: 7.01e-08

Can we really pool the data from the different beaches?

```
> library(lattice)  
> dotplot(residuals(mod1)~rikz$Beach,jitter.x=TRUE,pch=1)
```



```
> xyplot(ShannonW~NAP|Beach,data=rikz)
```




```
> mod2 <- lm(ShannonW~NAP+Beach,data=rikz)  
> summary(mod2)  
[. . ]  
                Estimate Std. Error t value Pr(>|t|)  
(Intercept) 0.620395    0.086771    7.150 2.45e-08 ***
```

NAP	-0.207945	0.031366	-6.630	1.15e-07	***
Beach2	0.004665	0.123148	0.038	0.96999	
Beach3	-0.328311	0.121791	-2.696	0.01072	*
Beach4	-0.345859	0.125102	-2.765	0.00903	**
Beach5	0.025370	0.125550	0.202	0.84103	
Beach6	-0.185777	0.122910	-1.511	0.13964	
Beach7	-0.363441	0.128055	-2.838	0.00750	**
Beach8	-0.125882	0.124720	-1.009	0.31975	
Beach9	-0.077617	0.125268	-0.620	0.53953	

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

Residual standard error: 0.1911 on 35 degrees of freedom  
Multiple R-squared: 0.7148, Adjusted R-squared: 0.6415  
F-statistic: 9.747 on 9 and 35 DF, p-value: 2.850e-07



```
> mod3 <- lm(ShannonW~NAP*Beach,data=rikz)
> summary(mod3)
[...]
```


	Estimate	Std. Error	t value	Pr(> t )	
(Intercept)	0.663388	0.097592	6.798	2.67e-07	***
NAP	-0.118188	0.113333	-1.043	0.30627	
Beach2	0.002042	0.129370	0.016	0.98752	
Beach3	-0.371311	0.126891	-2.926	0.00688	**
Beach4	-0.413504	0.135632	-3.049	0.00510	**
Beach5	0.130807	0.142539	0.918	0.36691	
Beach6	-0.245660	0.127976	-1.920	0.06554	.
Beach7	-0.528549	0.165405	-3.195	0.00354	**
Beach8	-0.188925	0.131537	-1.436	0.16241	
Beach9	-0.086618	0.136291	-0.636	0.53043	
NAP:Beach2	-0.236876	0.145877	-1.624	0.11604	
NAP:Beach3	-0.075049	0.149266	-0.503	0.61919	
NAP:Beach4	-0.044999	0.142857	-0.315	0.75519	
NAP:Beach5	-0.335176	0.156110	-2.147	0.04093	*
NAP:Beach6	-0.017519	0.133795	-0.131	0.89679	
NAP:Beach7	0.050573	0.166465	0.304	0.76361	



```

NAP:Beach8 -0.049864  0.132654 -0.376  0.70993
NAP:Beach9 -0.149265  0.142921 -1.044  0.30556
[. .]

```



```

> anova(mod1,mod2,mod3)
[. .]
Model 1: ShannonW ~ NAP
Model 2: ShannonW ~ NAP + Beach
Model 3: ShannonW ~ NAP * Beach
  Res.Df    RSS Df Sum of Sq    F Pr(>F)
1      43 2.26293
2      35 1.27812  8   0.98481 3.7433 0.00458 **
3      27 0.88792  8   0.39020 1.4831 0.20951
[. .]
> AIC(mod1,mod2,mod3)
      df      AIC
mod1  3 -0.8457103
mod2 11 -10.5528837
mod3 19 -10.9445763

```

- Model 2 (where beaches have an additive effect) is significantly better than model 1, and model 3 is not significantly better than model 2.
- AIC favours (takes the lowest value for) model 3.
- How can AIC be negative? It is still defined as  $-2 \cdot \log(L) + 2 \cdot d$ , whereas  $d$  is the number of parameters and  $L$  is the likelihood of the ML parameter estimate in the model. Since the response variable ShannonW has a continuous distribution (assumed to be normal), the likelihood is the probability density of the data given the parameter values. Densities can be larger than 1. If  $L > 1$ , then  $\log(L) > 0$  and AIC can be negative.
- We are not interested in the precise effect of each beach
- We do not want to estimate 8 extra parameters for the beaches

- Is there another way to take the difference between the beaches into account?
- Assume that the effect  $\alpha_k$  of beach  $k$  is random. Do not estimate all  $\alpha_k$  but only their standard deviation  $\sigma_\alpha$ .

### Mixed-effects model

Let  $S_i$  and  $N_i$  be the ShannonW and the NAP observed at plot  $i$ , which is on beach  $k_i$ .

$$S_i = a + b \cdot N_i + \alpha_{k_i} + \varepsilon_i$$

$\varepsilon_1, \varepsilon_2, \dots, \varepsilon_{45}$  are independently  $\mathcal{N}(0, \sigma^2)$ -distributed.  $\alpha_1, \alpha_2, \dots, \alpha_9$  are independently  $\mathcal{N}(0, \sigma_\alpha^2)$ -distributed. Mixed-effects:  $a$  and  $b$  are *deterministic*,  $\alpha_1, \alpha_2, \dots, \alpha_9$  are *random*. [0.5cm] To be estimated:  $a, b, \sigma_\alpha, \sigma$ .

```
> library(lme4)
> mmod0 <- lmer(ShannonW~1+NAP+(1|Beach),
+              data=rikz,verbose=TRUE)
0:    -3.0154190: 0.730297
1:    -3.0321775: 0.694973
2:    -3.0321844: 0.694215
3:    -3.0321844: 0.694257
4:    -3.0321844: 0.694257
```

It may be necessary to first install the lme4 package:

```
> install.packages("lme4")
```

This will only work if the R version installed is not too old.

```
> summary(mmod0)
Linear mixed model fit by REML
Formula: ShannonW ~ 1 + NAP + (1 | Beach)
Data: rikz
      AIC   BIC logLik deviance REMLdev
4.968 12.19  1.516  -12.27  -3.032
Random effects:
Groups   Name          Variance Std.Dev.
Beach    (Intercept) 0.017595 0.13264
Residual                    0.036504 0.19106
Number of obs: 45, groups: Beach, 9

Fixed effects:
              Estimate Std. Error t value
(Intercept)  0.46722    0.05366   8.707
NAP          -0.21380    0.03060  -6.987

Correlation of Fixed Effects:
      (Intr)
NAP -0.198
```

What is REML?

Why are there  
 $t$ -values but no  $p$ -  
values?

### REML vs. ML

- ML (Maximum Likelihood): estimate all parameters (here  $a, b, \sigma_\alpha, \sigma$ ) by maximizing their joint likelihood.
- REML (Restricted Maximum Likelihood): first estimate variance parameters (here  $\sigma_\alpha, \sigma$ ) from the components of the response space that are orthogonal on components that can be explained by fixed effects. Using these estimates, the coefficients of the fixed effects (here  $a$  and  $b$ ) are estimated with ML.

- Comparable to estimation of  $\sigma^2$  from sample  $X_1, \dots, X_n$  by  $\frac{1}{n-1} \sum_i (\mu_X - X_i)^2$  instead of the biased ML estimator  $\frac{1}{n} \sum_i (\mu_X - X_i)^2$
- Also for fitting parameters of mixed-effects models, ML estimation is biased and REML is usually preferred.
- ML estimation should be used when a likelihood ratio test shall be applied to models with different fixed effects and the same random effects.

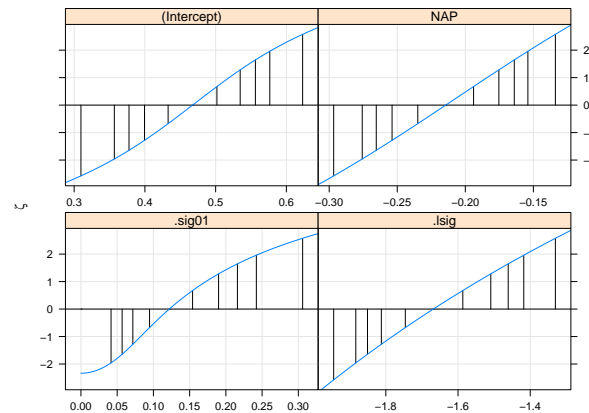
### Why no $p$ -values for the $t$ -values?

- The  $t$ -values computed like in the usual linear model, but in the case of mixed-effects models they are in general not  $t$ -distributed (under the null hypothesis). Thus, it is not clear how to get  $p$ -values from the  $t$ -values.
- Some other programs give  $p$ -values which can be very imprecise.
- Exception: small balanced datasets. Here,  $t$ -values are approximately  $t$ -distributed and  $|t| > 2$  usually indicates significance on the 5% level.

The computation of confidence ranges in lme4 is based on likelihood ratios. The profiles in the first plots are likelihood ratios.

Vertical lines indicate 50%, 80%, 90%, 95%, and 99% confidence intervals.

```
> pr0 <- profile(mmod0)
> xyplot(pr0)
```



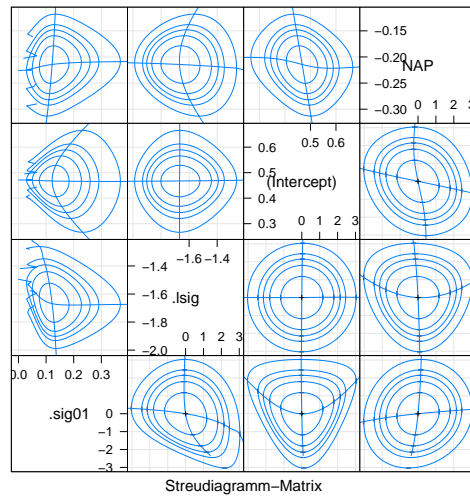
Display confidence intervals (per default 95%).

```
> confint(pr0)
                2.5 %      97.5 %
.sig01         0.0415353  0.2420610
.lsig          -1.8831563 -1.4178656
(Intercept)    0.3568131  0.5765609
NAP            -0.2757402 -0.1540307
```

The next plot shows pairwise confidence ranges (again 50%, 80%, 90%, 95%, and 99%).

Panels above the diagonal are in the original scale of the parameters, panels below the diagonal are on the  $\sqrt{\text{likelihoodratio}}$  scale.

```
> splom(pr0)
```



Useful information on methods for getting p values and confidence intervals are given in the R online help page for `pvalues` in the `lme4` package.

According to this, using likelihood profiles might be inaccurate when the number of groups is below 50 (like for the beaches), and using e.g. `PBmodcomp` package might be preferable in this case.

### 3 How to select a mixed-effects model

#### General considerations

- What is the purpose of the model?
  1. Making predictions as precise as possible
  2. or to understand what the most influential parameters are?
- In the first case AIC may be appropriate.
- In the second case it may be better to use likelihood-ratio tests and remove all parameters which do not significantly improve the fit.
- Variable selection should not only depend on statistics but also on the relevance of the parameter for the biological question.

We begin with a comparison of two nested models with the same random parameters

```
> B0 <- lmer(ShannonW~1+exposure*NAP+(1|Beach),
+           data=rikz,verbose=TRUE)
0:    -3.1079878: 0.730297
1:    -3.9300741: 0.431067
2:    -3.9406377: 0.471682
3:    -3.9424344: 0.460627
4:    -3.9424352: 0.460386
5:    -3.9424352: 0.460388
```

```

> summary(B0)
Linear mixed model fit by REML
Formula: ShannonW ~ 1 + exposure * NAP + (1 | Beach)
Data: rikz
AIC   BIC logLik deviance REMLdev
8.058 18.90 1.971  -23.54  -3.942
Random effects:
Groups   Name          Variance Std.Dev.
Beach    (Intercept) 0.0069857 0.08358
Residual                   0.0329581 0.18154
Number of obs: 45, groups: Beach, 9

```

```

Fixed effects:
              Estimate Std. Error t value
(Intercept)  1.92603    0.44765   4.303
exposure     -0.14294    0.04368  -3.272
[. .]

```

```

> summary(B0)
[. .]
Fixed effects:
              Estimate Std. Error t value
(Intercept)  1.92603    0.44765   4.303
exposure     -0.14294    0.04368  -3.272
NAP          -0.89673    0.33620  -2.667
exposure:NAP  0.06661    0.03270   2.037

```

```

Correlation of Fixed Effects:
              (Intr) exposr NAP
exposure     -0.996
NAP          -0.233  0.235
exposur:NAP  0.235 -0.238 -0.996

```

```

> B1 <- update(B0, ~.-exposure:NAP)
0:   -3.5739227: 0.730297
1:   -4.1112507: 0.00000
2:   -4.1112507: 3.65618e-06
3:   -4.1112511: 0.000144407
4:   -4.1650501: 0.0555603
5:   -4.7984144: 0.246592
6:   -5.0104872: 0.374551
7:   -5.0104967: 0.373651
8:   -5.0104968: 0.373725
9:   -5.0104968: 0.373722

```

```

> summary(B1)
Linear mixed model fit by REML
Formula: ShannonW ~ exposure + NAP + (1 | Beach)
Data: rikz
AIC   BIC logLik deviance REMLdev
4.99 14.02 2.505  -19.62  -5.01
Random effects:

```

Groups	Name	Variance	Std.Dev.
Beach	(Intercept)	0.0050958	0.071385
Residual		0.0364847	0.191010

Number of obs: 45, groups: Beach, 9


Fixed effects:

	Estimate	Std. Error	t value
(Intercept)	1.71143	0.41570	4.117
exposure	-0.12166	0.04053	-3.001
NAP	-0.21555	0.02990	-7.209


Correlation of Fixed Effects:

	(Intr)	exposr
exposure	-0.996	
NAP	0.015	-0.040

> dotplot(ranef(B1,postVar=TRUE))



> qqmath(ranef(B1,postVar=TRUE))



```

> anova(B0,B1)
Data: rikz
Models:
B1: ShannonW ~ exposure + NAP + (1 | Beach)
B0: ShannonW ~ 1 + exposure * NAP + (1 | Beach)
  Df      AIC      BIC  logLik  Chisq Chi Df Pr(>Chisq)
B1  5  -9.6222 -0.58886  9.8111
B0  6 -11.5446 -0.70461 11.7723 3.9224    1  0.04765 *
---
Signif. codes:  0 *** 0.001 ** 0.01 * 0.05 . 0.1  1

```

The  $p$ -value of 0.04765 may be imprecise because REML was used. To get a slightly more reliable  $p$ -value we fit both models with ML.

```

> B1ML <- update(B1,REML=FALSE)
0:  -16.875166: 0.730297
1:  -19.518802: 0.000000
2:  -19.518802: 3.18807e-06

```

```

> B0ML <- update(B0,REML=FALSE)
0:  -21.820395: 0.730297
1:  -22.739010: 0.000000
2:  -22.739010: 7.52231e-06
3:  -22.739011: 0.000259128
4:  -22.880262: 0.0868904
5:  -23.616152: 0.288666
6:  -23.703097: 0.366825
7:  -23.703443: 0.361997
8:  -23.703444: 0.362261
9:  -23.703444: 0.362262

```

```

> anova(B1ML,B0ML)
Data: rikz
Models:
B1ML: ShannonW ~ exposure + NAP + (1 | Beach)
B0ML: ShannonW ~ 1 + exposure * NAP + (1 | Beach)
  Df      AIC      BIC  logLik  Chisq Chi Df Pr(>Chisq)
B1ML 5  -9.5188 -0.48549  9.7594
B0ML 6 -11.7034 -0.86347 11.8517 4.1846    1  0.04079 *
---
Signif. codes:  0 *** 0.001 ** 0.01 * 0.05 . 0.1  1

```

The model B0 (or B0ML) with interaction between NAP and exposure fits significantly better.

Zuur *et al.*<sup>1</sup> recommend the following strategy for the selection of random and fixed parameters.

1. Start with a model that contains as many of the relevant parameters and interactions as possible.
2. First select random parameters. To decide between models which have different random parameters, fit models with REML and choose model of minimal AIC.

---

<sup>1</sup> A.F. Zuur, E.N. Ieno, N.J. Walker, A.A. Saveliev, G.M. Smith (2009) *Mixed effects models and extensions in ecology with R*. Springer.

3. Now select fixed parameters. This can be done with the help of AIC or with likelihood ratio tests. If likelihood ratio tests are used, apply ML to fit the models to the data.
4. Never remove covariates that are still involved in interactions.
5. Fit the final model with REML.

```
> B2 <- lmer(ShannonW~1+(exposure+salinity+NAP+humus)^2+(1|Beach),
+           .... [TRUNCATED]
0:    -5.8942807: 0.730297
1:    -5.8978355: 0.838579
2:    -5.9174068: 0.789454
3:    -5.9175376: 0.784845
4:    -5.9175385: 0.785156
5:    -5.9175385: 0.785156
```

```
> summary(B2)
Linear mixed model fit by REML
Formula: ShannonW ~ 1 + (exposure + salinity + NAP + humus)^2
+ (1 | Beach)
```

```
Data: rikz
AIC   BIC logLik deviance REMLdev
20.08 43.57 2.959  -37.55  -5.918
Random effects:
Groups   Name              Variance Std.Dev.
Beach    (Intercept) 0.015652 0.12511
Residual                    0.025390 0.15934
Number of obs: 45, groups: Beach, 9
```

```
Fixed effects:
              Estimate Std. Error t value
(Intercept)  40.32851   21.45346   1.880
exposure     -4.16161    2.12543  -1.958
salinity     -1.38750    0.74706  -1.857
NAP           0.65692    0.92037   0.714
humus        40.36916   16.33692   2.471
exposure:salinity 0.14522    0.07425   1.956
exposure:NAP    0.08357    0.03615   2.312
exposure:humus -2.06931    0.78148  -2.648
salinity:NAP   -0.06337    0.02639  -2.401
salinity:humus -0.67100    0.43169  -1.554
NAP:humus      1.01589    0.59247   1.715
[...]
```

```
[...]
Correlation of Fixed Effects:
              (Intr) exposr slnty NAP    humus  exspr:s ex:NAP exspr:h sl:NAP
exposure     -0.998
salinity     -0.999  0.998
NAP          -0.079  0.062  0.074
humus        0.117 -0.142 -0.129  0.152
exspr:slnty  0.996 -0.999 -0.998 -0.057  0.155
exposur:NAP  0.173 -0.165 -0.177 -0.661  0.067  0.167
exposur:hms -0.285  0.295  0.306 -0.008 -0.694 -0.315 -0.331
salinty:NAP  0.005  0.012  0.003 -0.920 -0.224 -0.020  0.316  0.189
salinty:hms  0.025  0.002 -0.022 -0.197 -0.888 -0.006  0.128  0.286  0.174
NAP:humus    0.149 -0.145 -0.158  0.109 -0.119  0.154  0.215 -0.149 -0.270
slnty:
exposure
salinity
NAP
```



```

humus
expr:slnty
exposur:NAP
exposur:hms
salinty:NAP
salinty:hms
NAP:humus    0.266

> B3 <- lmer(ShannonW~1+(exposure+salinity+NAP+humus)^2
+          +(1+NAP|Beach), data=rikz,verbose=TRUE)
 0:   -1.5323306: 0.730297 0.700275 0.00000
 1:   -5.4597824: 1.16363 0.00000 -0.193541
 2:   -5.7892458: 1.04081 4.21403e-06 -0.0735358
 3:   -5.9531543: 0.913041 6.58205e-06 -0.188264
[. . .]
12:  -6.0798268: 0.829521 0.00000 -0.109646
13:  -6.0798268: 0.829519 0.00000 -0.109651
14:  -6.0798268: 0.829519 0.00000 -0.109651

```

Next, we fit a model where there is not only a random intercept for every beach but also a random coefficient of NAP. Again, let  $S_i$  and  $N_i$  be the ShannonW and the NAP observed at plot  $i$ , which is on beach  $k$ . The model says

$$S_i = a + [\text{fixed effects terms}] + \alpha_k + \beta_k \cdot N_i + \varepsilon_i.$$

$\varepsilon_1, \dots, \varepsilon_{45}$  are independently  $\mathcal{N}(0, \sigma^2)$ -distributed, and also  $\left( \begin{smallmatrix} \alpha_1 \\ \beta_1 \end{smallmatrix} \right), \dots, \left( \begin{smallmatrix} \alpha_9 \\ \beta_9 \end{smallmatrix} \right) \sim \mathcal{N}\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_\alpha^2 & \rho_{\alpha\beta} \\ \rho_{\alpha\beta} & \sigma_\beta^2 \end{pmatrix}\right)$ .

Besides the fixed-effects coefficients we have to estimate the variances  $\sigma$ ,  $\sigma_\alpha$  and  $\sigma_\beta$ , and the covariance  $\rho_{\alpha\beta}$ .

The pairs of  $\alpha$  and  $\beta$  are independent in the sense that  $\left( \begin{smallmatrix} \alpha_i \\ \beta_i \end{smallmatrix} \right)$  is independent of  $\left( \begin{smallmatrix} \alpha_j \\ \beta_j \end{smallmatrix} \right)$  for  $i \neq j$  (and independent of all other coefficients).

If we also want independence between  $\alpha_i$  and  $\beta_i$ , that is  $\rho_{\alpha\beta}$  restricted to 0, we need to replace (1+NAP|Beach) in the R model formula by

(1+NAP|| Beach) or, equivalently,  
(1| Beach)+(0+NAP| Beach) or (1| Beach)+(NAP-1| Beach).

```

> summary(B3)
[. . .]
   AIC   BIC logLik deviance REMLdev
23.92 51.02   3.04  -37.69   -6.08
Random effects:
 Groups   Name      Variance Std.Dev. Corr
 Beach   (Intercept) 0.01730652 0.131554
         NAP         0.00020808 0.014425 -1.000
 Residual                0.02515108 0.158591
Number of obs: 45, groups: Beach, 9

Fixed effects:
      Estimate Std. Error t value
(Intercept)  40.93936   21.72964   1.884
exposure     -4.22455    2.15239  -1.963
salinity     -1.40881    0.75753  -1.860
NAP           0.67896    0.93389   0.727
humus        40.42412   16.30099   2.480
exposure:salinity 0.14742    0.07528   1.958

```

```

exposure:NAP      0.08480    0.03663    2.315
exposure:humus   -2.10698    0.78447   -2.686
salinity:NAP     -0.06457    0.02668   -2.421
salinity:humus  -0.66081    0.42975   -1.538
NAP:humus        0.98569    0.58792    1.677
[.]

```

```

> B4 <- lmer(ShannonW~1+(exposure+salinity+NAP+humus)^2
+           +(1+NAP+exposure|Beach),data=rikz,verbose=TRUE)
 0:  -1.6536097: 0.730297 0.700275 0.0711568 0.00000 0.00000 0.00000
 1:  -3.4551238: 0.700827 0.500322 0.00000 -0.0225207 -0.207903 -0.232659
 2:  -5.5427746: 0.630604 0.176807 0.00000 -0.00365199 -0.00670291 -0.156884
 3:  -5.7754683: 0.635022 0.175143 0.00000 -0.00439670 0.0278828 -0.162942
[.]
267: -6.2309242: 0.0242437 1.71415e-05 7.45081e-05 -3.81706 3.10877 0.0228216
268: -6.2309242: 0.0242437 1.71415e-05 7.45081e-05 -3.81706 3.10877 0.0228216
269: -6.2309242: 0.0242437 1.71415e-05 7.45081e-05 -3.81706 3.10877 0.0228216

```

```

> summary(B4)
[.]
      AIC   BIC logLik deviance REMLdev
29.77 62.29  3.115  -38.41  -6.231
Random effects:
Groups   Name          Variance  Std.Dev.  Corr
Beach    (Intercept)  1.4754e-05 0.0038411
         NAP        2.1496e-04 0.0146616 -1.000
         exposure  1.4259e-04 0.0119410 1.000 -1.000
Residual                2.5102e-02 0.1584364
Number of obs: 45, groups: Beach, 9

```

```

Fixed effects:
Estimate Std. Error t value
(Intercept) 42.77575 20.81776 2.055
exposure     -4.41472 2.06398 -2.139
salinity     -1.47464 0.72415 -2.036
NAP           0.67731 0.93297 0.726
humus        40.48334 16.27068 2.488
exposure:salinity 0.15425 0.07203 2.142
exposure:NAP   0.08472 0.03662 2.313
exposure:humus -2.11373 0.78433 -2.695
salinity:NAP  -0.06448 0.02663 -2.422
salinity:humus -0.66074 0.42862 -1.542
NAP:humus      0.98457 0.58726 1.677
[.]

```

```

> B5 <- lmer(ShannonW~1+(exposure+salinity+NAP+humus)^2
+           +(1+NAP+exposure|Beach)+(1|Week),data=rikz,verbose=TRUE)
 0:  -1.9949020: 0.730297 0.700275 0.0711568 0.00000 0.00000 0.00000 0.486864
 1:  -2.5896994: 0.696448 0.572039 0.00000 -0.0146206 -0.256429 -0.147424 0.501253
 2:  -5.3531758: 0.661521 0.525272 0.00000 -0.00132594 -0.106532 -0.107579 0.500492
[.]
62:  -9.1285182: 4.30055 1.60825e-07 0.00000 0.0323185 -0.0979504 1.69609 0.928692
63:  -9.1285182: 4.30055 1.49005e-08 5.79252e-07 0.0323201 -0.0979503 1.69609 0.928692
64:  -9.1285182: 4.30055 0.00000 0.00000 0.0323202 -0.0979503 1.69609 0.928692

```

```

> summary(B5)
Linear mixed model fit by REML
Formula: ShannonW ~ 1 + (exposure + salinity + NAP + humus)^2
          + (1 + NAP + exposure | Beach) + (1 | Week)
Data: rikz
      AIC   BIC logLik deviance REMLdev
28.87 63.2  4.564  -40.07  -9.129
Random effects:
Groups   Name          Variance  Std.Dev.  Corr

```

```

Beach (Intercept) 0.44901920 0.670089
      NAP         0.00046904 0.021657 1.000
      exposure   0.00430800 0.065635 -1.000 -1.000
Week (Intercept) 0.02093924 0.144704
Residual          0.02427827 0.155815
Number of obs: 45, groups: Beach, 9; Week, 4

```

Fixed effects:

```

Estimate Std. Error t value
(Intercept) 21.42334 15.49344 1.383
exposure -2.34537 1.53851 -1.524
salinity -0.74919 0.54787 -1.367
NAP 0.72633 0.93509 0.777
humus 38.75912 15.71392 2.466
exposure:salinity 0.08391 0.05452 1.539
exposure:NAP 0.08035 0.03615 2.223
exposure:humus -2.01279 0.75489 -2.666
salinity:NAP -0.06471 0.02681 -2.414
salinity:humus -0.63360 0.41503 -1.527
NAP:humus 0.99119 0.58288 1.700
NAP:humus 0.292
[...]
```

```
> anova(B2,B3,B4,B5)
```

Data: rikz

Models:

```

B2: ShannonW ~ 1 + (exposure + salinity + NAP + humus)^2
    + (1 | Beach)
B3: ShannonW ~ 1 + (exposure + salinity + NAP + humus)^2
    + (1 + NAP | Beach)
B4: ShannonW ~ 1 + (exposure + salinity + NAP + humus)^2
    + (1 + NAP + exposure | Beach)
B5: ShannonW ~ 1 + (exposure + salinity + NAP + humus)^2
    + (1 + NAP + exposure | Beach) + (1 | Week)
  Df    AIC    BIC logLik Chisq Chi Df Pr(>Chisq)
B2 13 -11.5528 11.934 18.776
B3 15 -7.6911 19.409 18.846 0.1383 2 0.9332
B4 18 -2.4079 30.112 19.204 0.7168 3 0.8693
B5 19 -2.0744 32.252 20.037 1.6665 1 0.1967

```

## How about models with less correlation in random effects?

```

> B5b <- lmer(ShannonW~1+(exposure+salinity+NAP+humus)^2+
              (1+NAP|Beach)+(exposure+0|Beach)+(1|Week),data=rikz)
> B5c <- lmer(ShannonW~1+(exposure+salinity+NAP+humus)^2+
              (1+NAP||Beach)+(exposure+0|Beach)+(1|Week),data=rikz)
> anova(B2,B5b,B5c)
refitting model(s) with ML (instead of REML)
Data: rikz
Models:
B2: ShannonW ~ 1 + (exposure + salinity + NAP + humus)^2 + (1 | Beach)
B5c: ShannonW ~ 1 + (exposure + salinity + NAP + humus)^2 + ((1 |
B5c: Beach) + (0 + NAP | Beach)) + (exposure + 0 | Beach) + (1 |
B5c: Week)
B5b: ShannonW ~ 1 + (exposure + salinity + NAP + humus)^2 + (1 + NAP |
B5b: Beach) + (exposure + 0 | Beach) + (1 | Week)
  Df    AIC    BIC logLik deviance Chisq Chi Df Pr(>Chisq)
B2 13 -12.2395 11.247 19.120 -38.240
B5c 16 -9.8246 19.082 20.912 -41.825 3.585 3 0.3099
B5b 17 -7.8246 22.889 20.912 -41.825 0.000 1 1.0000

```

Don't trust the  $p$ -values on the previous slides! One problem is that the models were fitted with REML. Another problem is that the null hypotheses (e.g.  $\sigma_\beta = 0$  in the case of B2/B3) are on the boundary of the parameter space.  $\sigma_\beta$  can only be  $\geq 0$ , and deviations from  $\sigma_\beta = 0$  are thus only possible in one direction.

The  $\chi^2$ -approximation of likelihood ratio tests are only reliable when deviations from the expectation under the null hypothesis are possible in all directions, for example if the null hypothesis  $\theta = 0$  is tested for some parameter  $\theta$ , and estimates of  $\theta$  can lead to positive as well as negative values.

Thus, we rather base our decision on the AIC values. This is, of course, also not stringent. However, in our case, all criteria favor model B2.

Now we keep the random effects “(1 | Beach)” and search for a good selection of fixed effects. Here, our aim is to get a simple, understandable model. Therefore apply a strict criterion and remove all variables which do not significantly improve the model fit on the 5% level.

(If our aim was a model to make good predictions, we would rather take AIC as a criterion, which leaves more parameters in.)

```
> summary(B2)
Linear mixed model fit by REML
Formula: ShannonW ~ 1 + (exposure + salinity + NAP + humus)^2
          + (1 | Beach)
[.]
Fixed effects:

```

	Estimate	Std. Error	t value
(Intercept)	40.32851	21.45346	1.880
exposure	-4.16161	2.12543	-1.958
salinity	-1.38750	0.74706	-1.857
NAP	0.65692	0.92037	0.714
humus	40.36916	16.33692	2.471
exposure:salinity	0.14522	0.07425	1.956
exposure:NAP	0.08357	0.03615	2.312
exposure:humus	-2.06931	0.78148	-2.648
salinity:NAP	-0.06337	0.02639	-2.401
salinity:humus	-0.67100	0.43169	-1.554
NAP:humus	1.01589	0.59247	1.715

```
[.]
> B6 <- update(B2, ~.-salinity:humus)
0:    -3.3840456: 0.730297
1:    -3.3848655: 0.748586
2:    -3.3851681: 0.742306
3:    -3.3851681: 0.742296
4:    -3.3851681: 0.742299
5:    -3.3851681: 0.742299

> anova(B6,B2)
Data: rikz
Models:
B6: ShannonW ~ exposure + salinity + NAP + humus + (1 | Beach) +
B6:    exposure:salinity + exposure:NAP + exposure:humus + salinity:NAP +
B6:    NAP:humus
B2: ShannonW ~ 1 + (exposure + salinity + NAP + humus)^2 + (1 | Beach)
Df    AIC    BIC logLik Chisq Chi Df Pr(>Chisq)
B6 12 -10.720 10.960 17.360
B2 13 -11.553 11.934 18.776 2.8329    1    0.09235 .
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Remember that we should better use ML instead of REML if we want to apply likelihood ratio tests!

```
> B2ML <- update(B2, REML=FALSE)
0:    -37.815253: 0.730297
1:    -38.227458: 0.520323
```

```

2:    -38.238951: 0.555035
3:    -38.239535: 0.548840
4:    -38.239535: 0.548745
5:    -38.239535: 0.548714
6:    -38.239535: 0.548724

> B6ML <- update(B6,REML=FALSE)
0:    -34.786862: 0.730297
1:    -35.481002: 0.455308
2:    -35.502068: 0.503911
3:    -35.503579: 0.494055
4:    -35.503579: 0.493868
5:    -35.503579: 0.493857
6:    -35.503579: 0.493844
7:    -35.503579: 0.493844

> anova(B6ML,B2ML)
Data: rikz
Models:
B6ML: ShannonW ~ exposure + salinity + NAP + humus + (1 | Beach) +
B6ML:    exposure:salinity + exposure:NAP + exposure:humus + salinity:NAP +
B6ML:    NAP:humus
B2ML: ShannonW ~ 1 + (exposure + salinity + NAP + humus)^2 + (1 | Beach)
      Df      AIC      BIC logLik Chisq Chi Df Pr(>Chisq)
B6ML 12 -11.504 10.176 17.752
B2ML 13 -12.239 11.247 19.120 2.736      1 0.09811 .
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

```

On Ben Bolker's GLMM FAQ:

“Consider *not* testing the significance of random effects. If the random effect is part of the experimental design, this procedure may be considered sacrificial pseudoreplication (Hurlbert 1984). Using stepwise approaches to eliminate non-significant terms in order to squeeze more significance out of the remaining terms is dangerous in any case.”

## References

- [1] Ben Bolker et al. (2020) GLMM FAQ <https://bbolker.github.io/mixedmodels-misc/glmmFAQ.html>
- [2] Hurlbert, S. (1984) Pseudoreplication and the Design of Ecological Field Experiments. *Ecological Monographs* **54**: 187-211.

### Further reading

## References

- [1] lme4 online help page on p values (online available e.g. on <https://rdrr.io/cran/lme4/man/pvalues.html>)
- [2] B.M. Bolker, M.E. Brooks, C.J. Clark, S.W. Geange, J.R. Poulsen, M.H.H. Stevens, J.-S.S. White (2009) Generalized linear mixed models: a practical guide for ecology and evolution *Trends in ecology & evolution* **24**:127-135

- [3] Allegue, H., Araya-Ajoy, Y.G., Dingemanse, N.J., Dochtermann N.A., Garamszegi, L.Z., Nakagawa, S., Rale, D., Schielzeth, H. and Westneat, D.F. (2016) SQuID – Statistical Quantification of Individual Differences: an educational and statistical tool for understanding multi-level phenotypic data in linear mixed models. *Methods in Ecology and Evolution* **8**:257-267.

**Some of the publications by Douglas Bates**

## References

- [1] D. Bates, M. Mächler, B.M. Bolker, S.C. Walker (2015) Fitting linear mixed-effects models using lme4 *Journal of Statistical Software* **67**: 1–48 <https://www.jstatsoft.org/v067/i01>
- [2] D. Bates, R. Kliegl, S. Vasishth, H. Baayen (2015) Parsimonious mixed models <https://arxiv.org/abs/1506.04967>
- [3] H. Matuschek, R. Kliegl, S. Vasishth, H. Baayen, D. Bates (2017) Balancing Type I error and power in linear mixed models *Journal of Memory and Language* **94**: 305–315 <http://dx.doi.org/10.1016/j.jml.2017.01.001>

**Some of what you should be able to explain**

- Assumptions of mixed-effects models
- Why random-effect factors are sometimes better than fixed-effect factors
- How to translate an lme4 mixed-effects model formula into a precise mathematical model specification and vice versa
- When to use REML and when using ML
- Parametric bootstrapping
- Possible approaches and problems of model selection, also with random-effect factors