Aufgabe 2

Betrachten Sie erneut den Datensatz zum Blutdruck. Als Signifikanzniveau wird in dieser Aufgabe $\alpha = 5\%$ verwendet.

(a) Fitten Sie folgendes Modell in R, wobei $\text{gender}_i = 0$, falls Person i weiblich und $\text{gender}_i = 1$ falls Person i männlich ist.

$$\begin{aligned}\text{SBD}_{ij} = \quad & \beta_1 \cdot I(\text{gender}_i = 0) + \beta_2 \cdot I(\text{gender}_i = 1) + \beta_3 \cdot \text{dosis}_{ij} \cdot I(\text{gender}_i = 0) + \\ & \beta_4 \cdot \text{dosis}_{ij} \cdot I(\text{gender}_i = 1) + b_i + \varepsilon_{ij},\end{aligned}$$

$$b_i \overset{\text{iid}}{\sim} N(0, \tau^2),\ \varepsilon_{ij} \overset{\text{iid}}{\sim} N(0, \sigma^2),\quad b_i,\ \varepsilon_{ij} \text{ unabhängig} \quad \forall\, i, j$$

(b) Testen Sie die Hypothesen $H_0 : \beta_3 = \beta_4 = 0$ und $H_0 : \beta_3 = \beta_4$. Was bedeuten diese Hypothesen inhaltlich?

(c) Testen Sie ob die zufälligen Intercepts im Modell enthalten sein sollen. Formulieren Sie dazu die Nullhypothese und führen Sie den Test approximativ und exakt (mit der Funktion `exactRLRT()` aus dem R-Paket `RLRsim`) durch.

(d) Fitten Sie ein Modell, das zusätzlich für jede Person eine zufällige Steigung enthält, wobei zufälliger Intercept und zufällige Steigung korreliert sein können. Führen Sie einen approximativen Test durch, ob die zufällige Steigung im Modell enthalten sein soll.

(e) Berechnen Sie das konditionale AIC (cAIC) für das Modell aus (a) und (d). Welches Modell wird laut cAIC bevorzugt?
Hinweis: Funktion `cAIC()` aus dem R-Paket `cAIC4`.

(f) *Zusatzaufgabe*: Führen Sie den Test für die zufällige Steigung mit der Funktion `exactRLRT()` durch. Beachten Sie, dass Random Intercepts und Slopes unabhängig sein müssen, damit `exactRLRT()` angewendet werden kann.

Lösung 2

(a)
```
# create dummy-variables -> makes testing easier
blutdruck$genderFemaleDosis <- 0
blutdruck$genderFemaleDosis[blutdruck$gender==0] <-
  (blutdruck$dosis)[blutdruck$gender==0]

blutdruck$genderMaleDosis <- 0
blutdruck$genderMaleDosis[blutdruck$gender==1] <-
  (blutdruck$dosis)[blutdruck$gender==1]

# treat gender as factor
blutdruck$gender <- factor(blutdruck$gender)

# Model with random intercept and different slope for male and female
fm1 <- lme(SBD ~ -1 + gender + genderFemaleDosis + genderMaleDosis,
           random=~1|ID, data=blutdruck)
summary(fm1)
```

```
Linear mixed-effects model fit by REML
 Data: blutdruck
       AIC      BIC    logLik
  1160.889 1179.912 -574.4445

Random effects:
 Formula: ~1 | ID
        (Intercept) Residual
StdDev:    5.179378  4.88174

Fixed effects: SBD ~ -1 + gender + genderFemaleDosis + genderMaleDosis
                       Value Std.Error  DF   t-value p-value
gender0           153.81288 1.5857629  28  96.99614       0
gender1           151.22683 1.5857629  28  95.36534       0
genderFemaleDosis  -0.15363 0.0094787 149 -16.20803       0
genderMaleDosis    -0.17898 0.0094787 149 -18.88289       0
 Correlation:
                  gendr0 gendr1 gndrFD
gender1            0.000
genderFemaleDosis -0.428  0.000
genderMaleDosis    0.000 -0.428  0.000

Standardized Within-Group Residuals:
       Min         Q1        Med         Q3        Max
-2.1531996 -0.5242295 -0.1017439  0.4420274  3.4061271

Number of Observations: 180
Number of Groups: 30
```

(b)

- $H_0 : \beta_3 = \beta_4 = 0$: Dosis hat keinen Einfluss unabhängig vom Geschlecht
- $H_0 : \beta_3 = \beta_4$: Effekt der Dosis ist für beide Geschlechter gleich
- Verwende Wald-Test weil LQ-Test nicht möglich (Likelihoods wegen REML-Schätzung nicht vergleichbar)
- Alternativ LRT mit Modellen, die via ML geschätzt wurden

```r
###### (2b)
# Calculate the p-value of the Wald-test for dosis-effects
hatbeta <- fixef(fm1)[3:4]
tw <- hatbeta %*% solve(vcov(fm1)[3:4,3:4]) %*% hatbeta
1 - pchisq(tw, df=2)
```

```
     [,1]
[1,]    0
```

```
# reject H0 on 5% level

# EXTRA: using anova() on one model gives F-test of effects for single variables
# see ?anova.lme
anova(fm1)
```

```
                  numDF denDF  F-value p-value
gender                2    28 9630.638  <.0001
genderFemaleDosis     1   149  262.700  <.0001
genderMaleDosis       1   149  356.563  <.0001
```

```
### do a likelihood-ratio test (LRT), models have to be fitted using ML
fm1ML1 <- lme(SBD ~ -1 + gender + genderFemaleDosis + genderMaleDosis,
              random=~1|ID, data=blutdruck, method="ML")
fm1ML0 <- lme(SBD ~ -1 + gender, random=~1|ID, data=blutdruck, method="ML")
# LRT for models fitted by ML
anova(fm1ML0, fm1ML1)
```

```
       Model df      AIC      BIC    logLik   Test  L.Ratio p-value
fm1ML0     1  4 1393.804 1406.576 -692.9020
fm1ML1     2  6 1150.961 1170.119 -569.4807 1 vs 2 246.8428  <.0001
```

```
## or LRT by hand
tLQT <- 2*fm1ML1$logLik - 2*fm1ML0$logLik
1 - pchisq(tLQT, df=2)
```

```
[1] 0
```

```
### try LQT for models fitted with REML
fm1REML0 <- lme(SBD ~ -1 + gender, random=~1|ID, data=blutdruck, method="REML")
anova(fm1, fm1REML0) ## gives a warning, as it ist not meaningful!
```

```
         Model df      AIC      BIC    logLik   Test  L.Ratio p-value
fm1          1  6 1160.889 1179.912 -574.4445
fm1REML0     2  4 1388.759 1401.486 -690.3797 1 vs 2 231.8702  <.0001
```

```
#### use package pbkrtest, for F-test and for parametric bootstrap
library(pbkrtest)
```

```
### (1) use F-test with approximated df
# An approximate F-test based on the Kenward-Roger approach.
# ?KRmodcomp

## you have to fit the models using lmer to use KRmodcomp()
library(lme4)
fm1lmer <- lmer(SBD ~ -1 + gender + genderFemaleDosis + genderMaleDosis + (1|ID),
                data=blutdruck)
fm0lmer <- lmer(SBD ~ -1 + gender + (1|ID), data=blutdruck)

# F-test with approximated df, based on Kenward-Roger approximation
KRmodcomp(fm1lmer, fm0lmer)

F-test with Kenward-Roger approximation; computing time: 0.31 sec.
large : SBD ~ -1 + gender + genderFemaleDosis + genderMaleDosis + (1 |
    ID)
small : SBD ~ -1 + gender + (1 | ID)
        stat    ndf    ddf F.scaling   p.value
Ftest 309.63   2.00 148.00         1 < 2.2e-16 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

### (2) use parametric bootstrap methods
# (takes about 1 minute)
PBmodcomp(fm1lmer, fm0lmer)

Parametric bootstrap test; time: 24.39 sec; samples: 1000 extremes: 0;
large : SBD ~ -1 + gender + genderFemaleDosis + genderMaleDosis + (1 |
    ID)
small : SBD ~ -1 + gender + (1 | ID)
         stat df   p.value
LRT    246.86  2 < 2.2e-16 ***
PBtest 246.86     0.000999 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

```
# Calculate the p-value of the Wald-test for the second hypothesis
L <- t(c(1,-1)) # matrix with one row
hatbeta <- fixef(fm1)[3:4]
```

```
tw2 <- t(hatbeta)%*%t(L) %*% solve(L%*%vcov(fm1)[3:4,3:4]%*%t(L)) %*% L%*%hatbeta
1 - pchisq(tw2, df=2)

          [,1]
[1,] 0.1671753

# do not reject H0 on 5% level
```

(c)

- Test auf Varianzkomponenten des Random Intercepts
- $H_0 : \text{var}(b_{0i}) = \tau^2 = 0$
- Exakter Test mit `RLRsim`
- Approximation unter $H_0 : T_{LQT} \overset{a}{\sim} 0.5 \cdot \chi^2_q + 0.5\chi^2_{q-1}$

```
library(RLRsim)
# ?exactRLRT
# it is enough to specify m, as only one variance component is tested
exactRLRT(m = fm1)

simulated finite sample distribution of RLRT.

(p-value based on 10000 simulated values)

data:
RLRT = 71.08, p-value < 2.2e-16

# reject H0 on 5% level



### Test statistic to approximate p-values
# model under H0
fm1Fix <- lm(SBD ~ -1 + gender + genderFemaleDosis + genderMaleDosis,
             data=blutdruck)
# approximate likelihood-ratio test
T_RLRTc <- 2 * logLik(fm1, REML = TRUE)[1] - 2 * logLik(fm1Fix, REML = TRUE)[1]
# approximate p-value with mixture of chi^2 distributions, q = 1
0.5 *(1-pchisq(T_RLRTc, 1)) + 0.5 *(1-pchisq(T_RLRTc, 0))

[1] 0

# reject H0 on 5% level
```

(d)
- wieder Approximation über Mischverteilung
- jetzt $q = 2$, da sowohl slope als auch Kovarianz zwischen Intercept und Slope auf 0 getestet werden

```
# Model with random intercept (m0, model under H0)
fm1 <- lme(SBD ~ -1 + gender + genderFemaleDosis + genderMaleDosis,
           random = ~ 1 | ID, data = blutdruck)
summary(fm1)

Linear mixed-effects model fit by REML
 Data: blutdruck
       AIC      BIC    logLik
  1160.889 1179.912 -574.4445

Random effects:
 Formula: ~1 | ID
        (Intercept) Residual
StdDev:    5.179378  4.88174

Fixed effects: SBD ~ -1 + gender + genderFemaleDosis + genderMaleDosis
                      Value Std.Error  DF   t-value p-value
gender0           153.81288 1.5857629  28  96.99614       0
gender1           151.22683 1.5857629  28  95.36534       0
genderFemaleDosis  -0.15363 0.0094787 149 -16.20803       0
genderMaleDosis    -0.17898 0.0094787 149 -18.88289       0
 Correlation:
                  gendr0 gendr1 gndrFD
gender1            0.000
genderFemaleDosis -0.428  0.000
genderMaleDosis    0.000 -0.428  0.000

Standardized Within-Group Residuals:
       Min         Q1        Med         Q3        Max
-2.1531996 -0.5242295 -0.1017439  0.4420274  3.4061271

Number of Observations: 180
Number of Groups: 30

# Model with random intercept and random slope (mA, model under alternative)
fm2 <- lme(SBD ~ -1 + gender + genderFemaleDosis + genderMaleDosis,
           random = ~1 + dosis| ID, data = blutdruck)
summary(fm2)
```

```
Linear mixed-effects model fit by REML
 Data: blutdruck
      AIC      BIC    logLik
  960.547 985.9109 -472.2735

Random effects:
 Formula: ~1 + dosis | ID
 Structure: General positive-definite, Log-Cholesky parametrization
            StdDev     Corr
(Intercept) 0.17578119 (Intr)
dosis       0.07395999 0.939
Residual    2.35185842

Fixed effects: SBD ~ -1 + gender + genderFemaleDosis + genderMaleDosis
                      Value Std.Error  DF  t-value p-value
gender0           153.81288 0.4130630  28 372.3715       0
gender1           151.22683 0.4130630  28 366.1108       0
genderFemaleDosis  -0.15363 0.0196348 149  -7.8244       0
genderMaleDosis    -0.17898 0.0196348 149  -9.1157       0
 Correlation:
                  gendr0 gendr1 gndrFD
gender1            0.000
genderFemaleDosis -0.084  0.000
genderMaleDosis    0.000 -0.084  0.000

Standardized Within-Group Residuals:
       Min         Q1        Med         Q3        Max
-2.1340429 -0.6465530 -0.0215929  0.6304812  2.5156456

Number of Observations: 180
Number of Groups: 30
```

```
# Test statistic to approximate p-values
T_RLRTd <- 2 * logLik(fm2, REML = TRUE)[1] - 2 * logLik(fm1, REML = TRUE)[1]

0.5 *(1-pchisq(T_RLRTd, 2)) + 0.5 *(1-pchisq(T_RLRTd, 1))

[1] 0

# reject H0 on 5% level
```

(e) Modellvergleich bei Modellen mit verschiedenen Random Effects auch über das konditionale AIC möglich (cAIC).

```
library(lme4)
fm1lmer <- lmer(SBD ~ -1 + gender + genderFemaleDosis +
                  genderMaleDosis + (1|ID), data=blutdruck)
fm2lmer <- lmer(SBD ~ -1 + gender + genderFemaleDosis +
                  genderMaleDosis + (1 + dosis|ID), data=blutdruck)

# compute the cAIC
library(cAIC4)
# ?cAIC
c1 <- cAIC(fm1lmer)
c2 <- cAIC(fm2lmer)

## preferred model is the one with the minimum AIC value
c1$caic

[1] 1112.552

c2$caic

[1] 858.7664
```

(f)

```
## create groupedData-Object
blutdruck2 <- groupedData(SBD ~ dosis | ID, data = blutdruck,
                          labels = list(x = "Dosis",
                                        y = "Systolischer Blutdruck"),
                          units = list(x = "[mg]",
                                       y = "[mmHg]"))

# Model with random intercept and random slope (mA, model under alternative)
# random intercept and slope are independent!
fm2S <- lme(SBD ~ -1 + gender + genderFemaleDosis + genderMaleDosis,
                random=pdDiag(~dosis), data=blutdruck2)
summary(fm2S)

Linear mixed-effects model fit by REML
 Data: blutdruck2
       AIC      BIC    logLik
  958.8087 981.0021 -472.4044
```

```
Random effects:
 Formula: ~dosis | ID
 Structure: Diagonal
        (Intercept)      dosis Residual
StdDev:    0.369784 0.07529139 2.342923

Fixed effects: SBD ~ -1 + gender + genderFemaleDosis + genderMaleDosis
                      Value Std.Error  DF  t-value p-value
gender0           153.81288 0.4199985  28 366.2225       0
gender1           151.22683 0.4199985  28 360.0652       0
genderFemaleDosis  -0.15363 0.0199653 149  -7.6949       0
genderMaleDosis    -0.17898 0.0199653 149  -8.9648       0
 Correlation:
                  gendr0 gendr1 gndrFD
gender1            0.000
genderFemaleDosis -0.177  0.000
genderMaleDosis    0.000 -0.177  0.000

Standardized Within-Group Residuals:
        Min          Q1         Med          Q3         Max
-2.19730706 -0.63848623 -0.06270077  0.65784472  2.49701433

Number of Observations: 180
Number of Groups: 30

# Model with random slopes only
# (m, model containig only the random effect to be tested)
fm3S <- lme(SBD ~ -1 + gender + genderFemaleDosis + genderMaleDosis,
                random=~-1+dosis|ID, data=blutdruck2)
summary(fm3S)

Linear mixed-effects model fit by REML
 Data: blutdruck2
       AIC     BIC   logLik
  956.8441 975.867 -472.422

Random effects:
 Formula: ~-1 + dosis | ID
             dosis Residual
StdDev: 0.07542387 2.354575
```

```
Fixed effects: SBD ~ -1 + gender + genderFemaleDosis + genderMaleDosis
                      Value Std.Error  DF  t-value p-value
gender0           153.81288 0.4110362  28 374.2077       0
gender1           151.22683 0.4110362  28 367.9161       0
genderFemaleDosis  -0.15363 0.0200038 149  -7.6801       0
genderMaleDosis    -0.17898 0.0200038 149  -8.9475       0
 Correlation:
                  gendr0 gendr1 gndrFD
gender1            0.000
genderFemaleDosis -0.182  0.000
genderMaleDosis    0.000 -0.182  0.000

Standardized Within-Group Residuals:
        Min          Q1         Med          Q3         Max
-2.21468937 -0.65074389 -0.06219758  0.62989176  2.53727162

Number of Observations: 180
Number of Groups: 30

# Calcuate test on the random slopes
# simulate asymptotic p-value
exactRLRT(m = fm3S, mA = fm2S, m0 =fm1)

simulated finite sample distribution of RLRT.

(p-value based on 10000 simulated values)

data:
RLRT = 204.08, p-value < 2.2e-16

# reject H0 on 5% level

# and approximated
T_RLRTs <- 2 * logLik(fm2S, REML = TRUE)[1] - 2 * logLik(fm1, REML = TRUE)[1]
# approximate p-value with mixture of chi^2 distributions, q=1
0.5 *(1-pchisq(T_RLRTs, 1)) + 0.5 *(1-pchisq(T_RLRTs, 0))

[1] 0

# reject H0 on 5% level
```