

Data, data documentation and analysis scripts for
Determinants of English accents.

Martijn Wieling⁽¹⁾ & Jelke Bloem⁽²⁾ & R. Harald Baayen^(4,5) & John Nerbonne^(1,6)

¹University of Groningen, the Netherlands & ²University of Amsterdam, the Netherlands &
³Utrecht University, the Netherlands & ⁴Eberhard Karls University, Germany & ⁵University of Alberta, Canada & ⁶University of Freiburg, Germany

Status: Accepted to appear in *Proceedings of QITL-6*
Preprint: <http://www.martijnwieling.nl/files/WielingEtAl-accents-submitted.pdf>

Abstract

In this study we investigate which factors affect the degree of non-native accent of L2 speakers of English who learned English in school and mostly lived for some time in an anglophone setting. We use data from the Speech Accent Archive containing over 700 speakers speaking almost 160 different native languages. We show that besides several important predictors, including the age of English onset and length of anglophone residence, the linguistic distance between the speaker's native language and English is a significant predictor of the degree of non-native accent in pronunciation. This study extends an earlier study (Schepens et al., 2013) which only focused on Indo-European L2 learners of Dutch and used a general speaking proficiency measure.

Keywords: Second language learning, Mixed-effects regression, Pronunciation nativelikeness, Linguistic distances.

1 Packages and functions

```
library(mgcv)
library(itsadug)
library(car)

R.Version()$version.string

## [1] "R version 3.2.2 (2015-08-14)"

packageVersion('mgcv')

## [1] '1.8.8'

packageVersion('itsadug')

## [1] '1.0.1'

packageVersion('car')

## [1] '2.1.0'
```

2 English accents data set

```
load("data/accents.rda")

str(accents)

## 'data.frame': 712 obs. of  26 variables:
## $ Speaker      : Factor w/ 712 levels "afrikaans1","afrikaans2",...
## $ Language     : Factor w/ 159 levels "afrikaans","agni",...
## $ Iso          : Factor w/ 150 levels "afb","afr","aka",...
## $ Country      : Factor w/ 132 levels "afghanistan",...
## $ LD.log       : num -4.62 -4.98 -4.84 -4.43 -4.98 ...
## $ Nativelikeness : num NA NA NA 1.69 2.93 ...
## $ IsIESpeaker  : logi FALSE FALSE FALSE FALSE FALSE ...
## $ LDND         : num 99.6 99.6 99.6 99.6 99.6 99.6 99.6 99.6 ...
## $ LDN          : num 0.00939 0.00939 0.00939 0.00939 0.00939 ...
## $ PMI          : num 0.827 0.827 0.827 0.827 0.827 ...
## $ Age          : num 38 26 25 32 35 18 22 36 23 30 ...
## $ IsMale        : num 0 1 1 1 1 1 1 0 0 ...
## $ AEO          : num 12 5 15 11 15 6 16 12 10 14 ...
## $ LR           : num 4 2 1 8 0 0 0 1 0 4 ...
## $ NrLang        : int 0 2 2 0 1 1 2 2 2 1 ...
## $ GNI.log      : num 10.06 8.57 9.48 8.06 8.34 ...
## $ CountryEduYrs : num 7.78 6.4 7.92 5.56 4.37 ...
## $ LDND.c       : num 5.87 5.87 5.87 5.87 5.87 ...
## $ LDN.c        : num 0.000274 0.000274 0.000274 0.000274 0.000274 ...
## $ PMI.c        : num 0.0438 0.0438 0.0438 0.0438 0.0438 ...
## $ Age.c        : num 5.543 -6.457 -7.457 -0.457 2.543 ...
## $ AEO.c        : num -0.173 -7.173 2.827 -1.173 2.827 ...
## $ LR.c         : num -2.82 -4.82 -5.82 1.18 -6.82 ...
## $ NrLang.c     : num -1.336 0.664 0.664 -1.336 -0.336 ...
## $ GNI.log.c    : num 1.099 -0.386 0.523 -0.892 -0.614 ...
## $ CountryEduYrs.c: num -0.2266 -1.6013 -0.0868 -2.4422 -3.6353 ...
```

Legenda **accents**:

1. Speaker : the speaker
2. Language : the native language of the speaker
3. Iso : iso code of the language (as used in ASJP)
4. Country : the country of birth of the speaker
5. LD.log : log-transformed average Levenshtein distance with respect to average native American English (see Wieling, Bloem, Mignella et al., 2014)
6. Nativelikeness : the human-rated nativelikeness of the speaker (only available for 237 speakers; also from Wieling, Bloem, Mignella et al., 2014)
7. IsIESpeaker : if the speaker has an Indo-European Language (TRUE) or not (FALSE)

8. LDND : LDND (Wichmann et al.,2010) distance between English and the speaker's native language (on the basis of the 40-item ASJP set)
9. LDN : Normalized Levenshtein distance between English and the speaker's native language (on the basis of the 40-item ASJP set)
10. PMI : PMI-weighted pronunciation distance (using the method described by Jäger, 2013) between English and the speaker's native language (on the basis of the 40-item ASJP set)
11. Age : the age of the speaker
12. IsMale : the gender of the speaker (1: male, 0: female)
13. AEO : the age of English onset of the speaker
14. LR : the cumulative length of residence in an English-speaking country of a speaker
15. NrLang : the number of additional languages the speaker speaks (besides English)
16. GNI.log : the log-transformed gross national income of the speaker's country of birth (in 2011)
17. CountryEduYrs : the average number of years of education in the speaker's country (in 2011)
18. *.c: centered numerical predictors

3 Analysis and results

3.1 Descriptives

```
# gender distribution of participants





```

```

## [1] 0.8363424

# correlation between Nativelikeness ratings and Levenshtein distances (LD)
# (note the correlation might be slightly different from the one reported
# by Wieling, Bloem, Mignella et al. (2014) as this study uses only
# non-native speakers who learned English academically)
cor(accents$Nativelikeness, accents$LD.log, use='pairwise')

## [1] -0.7636405

# correlation between LD-based speaker ratings (LD.log) and country distances (LDND)
cor(accents$LD.log, accents$LDND, use='pairwise')

## [1] 0.2994741

# correlation between speaker ratings (Nativelikeness) and country distances (LDND)
cor(accents$Nativelikeness, accents$LDND, use='pairwise')

## [1] -0.3704642

# correlation between LDND and LDN
cor(accents$LDND, accents$LDN, use='pairwise')

## [1] 0.9742606

# correlation between LDND and PMI
cor(accents$LDND, accents$PMI, use='pairwise')

## [1] 0.9806162

# correlation between LDN and PMI
cor(accents$LDN, accents$PMI, use='pairwise')

## [1] 0.9622538

# number of speakers with Nativelikeness ratings
sum(!is.na(accents$Nativelikeness))

## [1] 237

```

3.2 Optimal model

```

# optimal model specification determined via model comparison (not shown)
m0 <- gam(LD.log ~ LDND.c + AEO.c + LR.c*Age.c+ CountryEduYrs.c + NrLang.c +
           s(Country, bs='re') + s(Language, bs='re'), data=accents,
           method='REML')

summary(m0) # random intercept for Language not needed

```

```

## 
## Family: gaussian
## Link function: identity
##
## Formula:
## LD.log ~ LDND.c + AEO.c + LR.c * Age.c + CountryEduYrs.c + NrLang.c +
##      s(Country, bs = "re") + s(Language, bs = "re")
##
## Parametric coefficients:
##             Estimate Std. Error t value Pr(>|t|)
## (Intercept) -4.815e+00 8.594e-03 -560.259 < 2e-16
## LDND.c       3.653e-03 9.102e-04   4.014 6.65e-05
## AEO.c        9.987e-03 1.033e-03   9.668 < 2e-16
## LR.c         -5.521e-03 1.046e-03  -5.279 1.76e-07
## Age.c        2.260e-03 6.899e-04   3.276 0.001108
## CountryEduYrs.c -1.964e-02 2.979e-03  -6.593 8.70e-11
## NrLang.c     -1.997e-02 5.833e-03  -3.424 0.000656
## LR.c:Age.c    1.046e-04 4.081e-05   2.564 0.010566
##
## Approximate significance of smooth terms:
##          edf Ref.df    F p-value
## s(Country) 3.193e+01    130 0.408 0.000359
## s(Language) 9.661e-04    157 0.000 0.561033
##
## R-sq.(adj) =  0.355 Deviance explained = 39.1%
## -REML = -221.6 Scale est. = 0.02621 n = 712

# final model is fit with ML, as in the GAM framework it offers
# more conservative p-values
m1 <- gam(LD.log ~ LDND.c + AEO.c + LR.c*Age.c+ CountryEduYrs.c + NrLang.c +
           s(Country,bs='re'), data=accents, method='ML')
summary(m1)

## 
## Family: gaussian
## Link function: identity
##
## Formula:
## LD.log ~ LDND.c + AEO.c + LR.c * Age.c + CountryEduYrs.c + NrLang.c +
##      s(Country, bs = "re")
##
## Parametric coefficients:
##             Estimate Std. Error t value Pr(>|t|)
## (Intercept) -4.815e+00 8.472e-03 -568.322 < 2e-16
## LDND.c       3.650e-03 9.011e-04   4.051 5.70e-05
## AEO.c        9.991e-03 1.032e-03   9.684 < 2e-16
## LR.c         -5.496e-03 1.044e-03  -5.265 1.89e-07
## Age.c        2.266e-03 6.894e-04   3.287 0.001065
## CountryEduYrs.c -1.958e-02 2.936e-03  -6.668 5.41e-11
## NrLang.c     -2.009e-02 5.824e-03  -3.450 0.000595
## LR.c:Age.c    1.037e-04 4.078e-05   2.542 0.011230

```

```

## 
## Approximate significance of smooth terms:
##          edf Ref.df      F  p-value
## s(Country) 29.87     130 0.382 0.000388
## 
## R-sq.(adj) =  0.353   Deviance explained = 38.7%
## -ML = -268.92   Scale est. = 0.026295 n = 712

# PMI measure as opposed to LDND measure is not better:
m1pmi <- gam(LD.log ~ PMI.c + AEO.c + LR.c*Age.c+ CountryEduYrs.c + NrLang.c +
               s(Country,bs='re'), data=accents, method='ML')
compareML(m1,m1pmi)

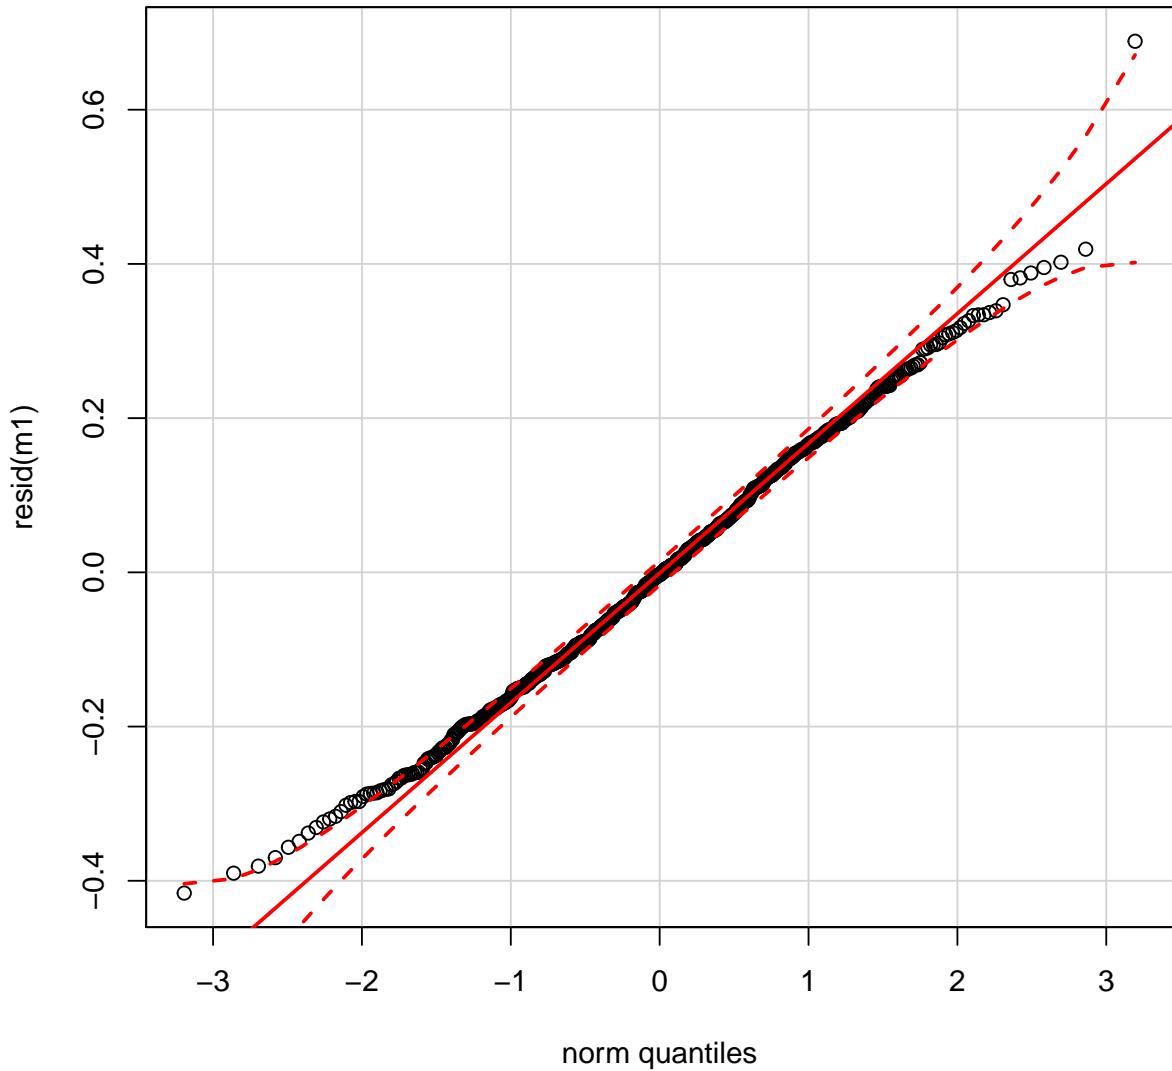
## m1: LD.log ~ LDND.c + AEO.c + LR.c * Age.c + CountryEduYrs.c + NrLang.c +
##       s(Country, bs = "re")
##
## m1pmi: LD.log ~ PMI.c + AEO.c + LR.c * Age.c + CountryEduYrs.c + NrLang.c +
##       s(Country, bs = "re")
##
## Model m1 preferred: lower ML score (0.366), and equal df (0.000).
## -----
##   Model      Score Edf Difference    Df
## 1 m1pmi -268.5529    9
## 2     m1 -268.9192    9      0.366 0.000
##
## AIC difference: -1.33, model m1 has lower AIC.

# LDN measure as opposed to LDND measure is not better:
m1ldn <- gam(LD.log ~ LDN.c + AEO.c + LR.c*Age.c+ CountryEduYrs.c + NrLang.c +
               s(Country,bs='re'), data=accents, method='ML')
compareML(m1,m1ldn)

## m1: LD.log ~ LDND.c + AEO.c + LR.c * Age.c + CountryEduYrs.c + NrLang.c +
##       s(Country, bs = "re")
##
## m1ldn: LD.log ~ LDN.c + AEO.c + LR.c * Age.c + CountryEduYrs.c + NrLang.c +
##       s(Country, bs = "re")
##
## Model m1 preferred: lower ML score (0.369), and equal df (0.000).
## -----
##   Model      Score Edf Difference    Df
## 1 m1ldn -268.5497    9
## 2     m1 -268.9192    9      0.369 0.000
##
## AIC difference: -1.33, model m1 has lower AIC.

qqp(resid(m1))

```



```

# 1 outlier at the top end is removed
accents2 = accents[resid(m1) < max(resid(m1)),]

# refitting without outlier
m1 <- gam(LD.log ~ LDND.c + AEO.c + LR.c*Age.c + CountryEduYrs.c + NrLang.c +
           s(Country,bs='re'), data=accents2, method='ML')

summary(m1) # summary of final model

##
## Family: gaussian

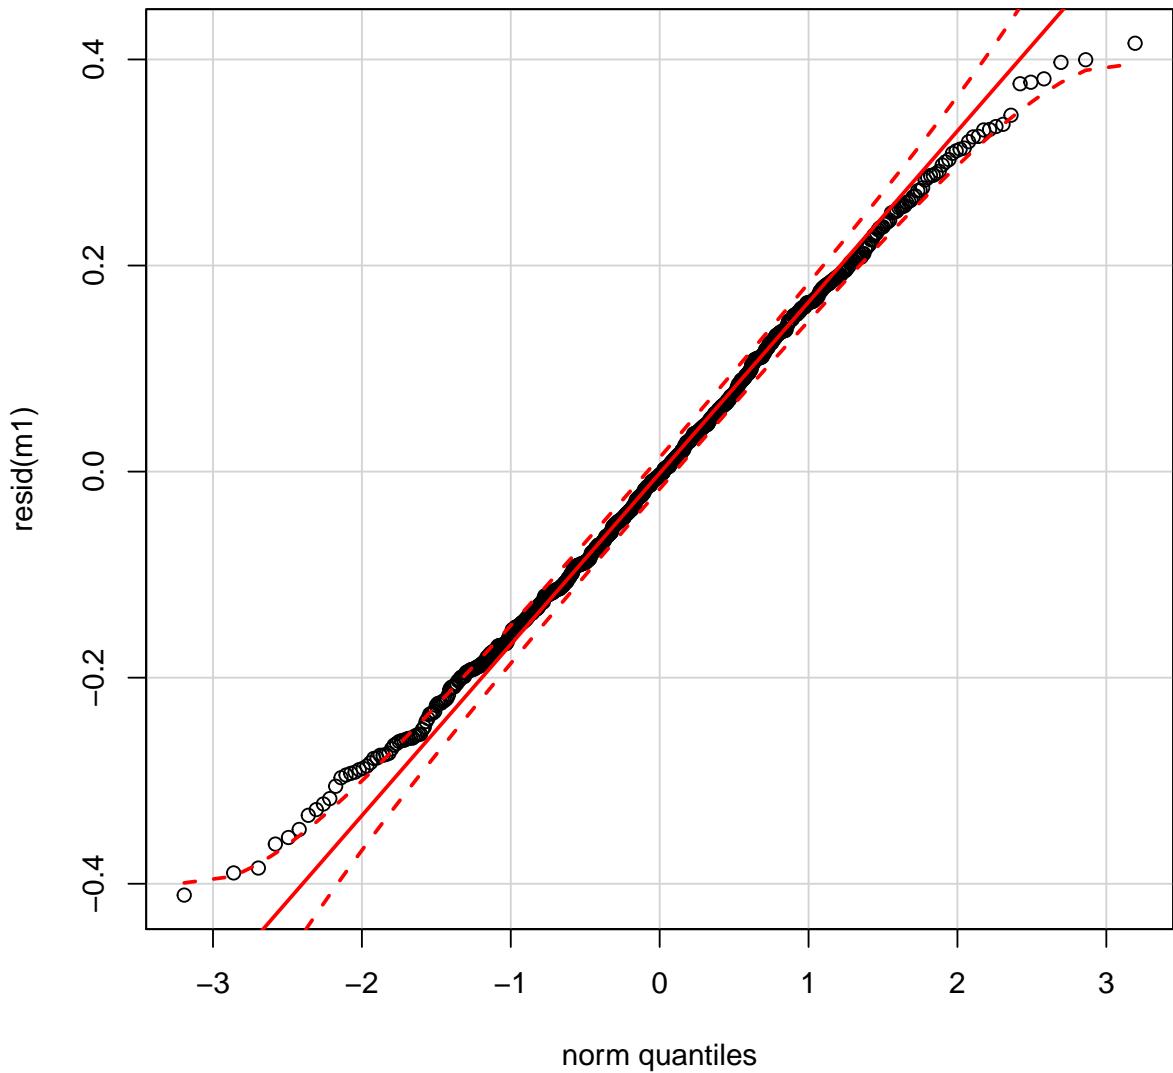
```

```

## Link function: identity
##
## Formula:
## LD.log ~ LDND.c + AEO.c + LR.c * Age.c + CountryEduYrs.c + NrLang.c +
##      s(Country, bs = "re")
##
## Parametric coefficients:
##                               Estimate Std. Error t value Pr(>|t|)
## (Intercept)           -4.816e+00 8.693e-03 -553.940 < 2e-16
## LDND.c                3.653e-03 9.134e-04   4.000 7.06e-05
## AEO.c                 9.926e-03 1.020e-03   9.732 < 2e-16
## LR.c                  -5.492e-03 1.034e-03  -5.313 1.47e-07
## Age.c                 2.143e-03 6.808e-04   3.147 0.001723
## CountryEduYrs.c     -1.971e-02 3.013e-03  -6.539 1.23e-10
## NrLang.c              -1.936e-02 5.763e-03  -3.359 0.000826
## LR.c:Age.c            1.077e-04 4.025e-05   2.675 0.007659
##
## Approximate significance of smooth terms:
##             edf Ref.df    F p-value
## s(Country) 35.67    130 0.505 1.55e-05
##
## R-sq.(adj) =  0.368 Deviance explained = 40.6%
## -ML = -277.77  Scale est. = 0.025279 n = 711

qqp(resid(m1)) # residuals approximately normally distributed

```



```
# effect sizes:
coef(m1)[["LDND.c"]]*(max(accents2$LDND.c)-min(accents2$LDND.c))

##      LDND.c
## 0.2073974

coef(m1)[["AE0.c"]]*(max(accents2$AE0.c)-min(accents2$AE0.c))

##      AE0.c
## 0.4466788

coef(m1)[["CountryEduYrs.c"]]*(max(accents2$CountryEduYrs.c)-min(accents2$CountryEduYrs.c))
```

```

## CountryEduYrs.c
## -0.2242543

coef(m1)[["NrLang.c"]]*(max(accents2$NrLang.c)-min(accents2$NrLang.c))

## NrLang.c
## -0.09679126

mx = -Inf
mn = Inf
accents2$LR.c_Age.c = accents2$LR.c * accents2$Age.c
fx = coef(m1)
for (i in 1:nrow(accents2)) {
  val = fx["LR.c"]*accents2[i,"LR.c"] +
    fx["Age.c"]*accents2[i,"Age.c"] +
    fx["LR.c:Age.c"]*accents2[i,"LR.c_Age.c"]
  names(val) = "LR.c*Age.c"
  if (val < mn) {
    mn = val
  }
  if (val > mx) {
    mx = val
  }
}
round(mx - mn,3) # 0.24

## LR.c*Age.c
## 0.24

# explained variance without random intercept for country
summary(tmp <- gam(LD.log ~ LDND.c + AEO.c + LR.c*Age.c+ CountryEduYrs.c + NrLang.c,
  data=accents2, method='ML'))

##
## Family: gaussian
## Link function: identity
##
## Formula:
## LD.log ~ LDND.c + AEO.c + LR.c * Age.c + CountryEduYrs.c + NrLang.c
##
## Parametric coefficients:
##                               Estimate Std. Error t value Pr(>|t|)
## (Intercept)      -4.813e+00  6.964e-03 -691.144 < 2e-16
## LDND.c          3.736e-03  7.828e-04   4.772 2.22e-06
## AEO.c           1.008e-02  1.012e-03   9.962 < 2e-16
## LR.c            -4.992e-03  1.018e-03  -4.906 1.15e-06
## Age.c           2.295e-03  6.843e-04   3.355 0.000838
## CountryEduYrs.c -1.809e-02  2.404e-03  -7.525 1.62e-13
## NrLang.c        -2.218e-02  5.690e-03  -3.898 0.000106
## LR.c:Age.c      8.907e-05  4.057e-05   2.195 0.028462
##

```

```

## 
## R-sq.(adj) =  0.308   Deviance explained = 31.4%
## -ML = -270.39   Scale est. = 0.027677  n = 711

# explained variance of LDND.c
summary(tmp <- gam(LD.log ~ LDND.c, data=accents2, method='ML'))

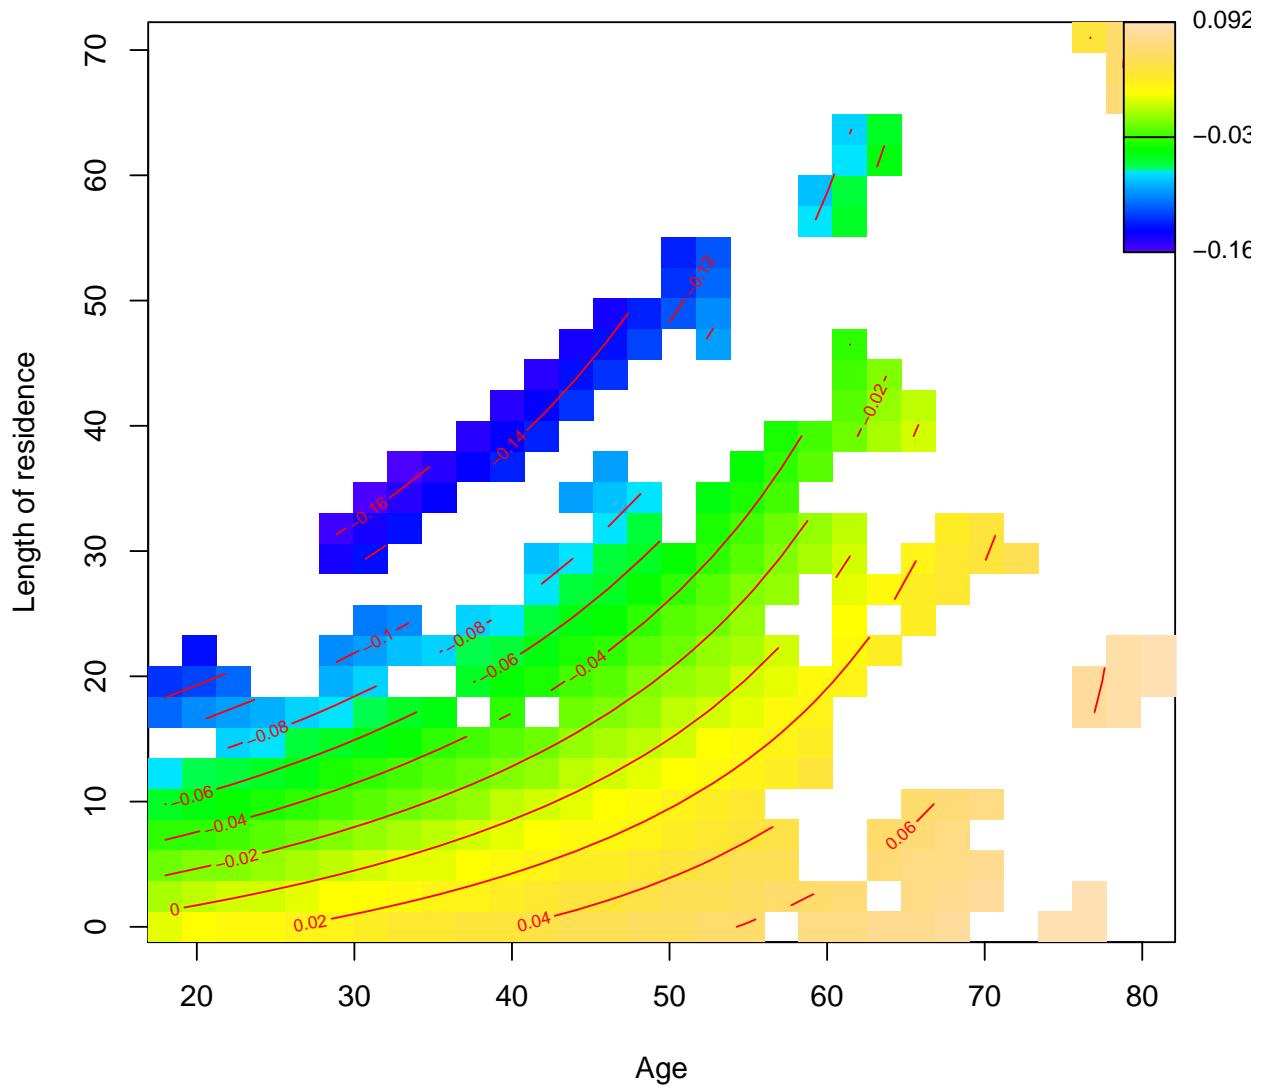
## 
## Family: gaussian
## Link function: identity
##
## Formula:
## LD.log ~ LDND.c
##
## Parametric coefficients:
##             Estimate Std. Error t value Pr(>|t|)    
## (Intercept) -4.8067695  0.0071551 -671.797  <2e-16  
## LDND.c       0.0065321  0.0007765    8.412  <2e-16  
## 
## 
## R-sq.(adj) =  0.0895   Deviance explained = 9.08%
## -ML = -169.98   Scale est. = 0.0364    n = 711

# visualization of interaction (refitting with tensor,
# which in this case simply models a linear interaction
# as the number of DF equals 3)
m1b <- gam(LD.log ~ LDND + AEO + te(Age,LR)+ CountryEduYrs + NrLang +
            s(Country,bs='re'), data=accents2, method='ML')

pvisgam(m1b, select=1, view=c('Age','LR'), too.far=0.04, hide.label=T,
         ylab='Length of residence', xlab='Age', main='')

## [1] "Tensor(s) to be plotted: te(Age,LR)"

```



3.3 Validation with nativelikeness ratings

```

nat = droplevels(accents[!is.na(accents$Nativelikeness),])
str(nat)

## 'data.frame': 237 obs. of  26 variables:
## $ Speaker      : Factor w/ 237 levels "afrikaans1","afrikaans2",...
## $ Language     : Factor w/ 87 levels "afrikaans","akan",...
## $ Iso          : Factor w/ 86 levels "afb","afr","aka",...
## $ Country      : Factor w/ 82 levels "afghanistan",...

```

```

## $ LD.log      : num -4.43 -4.98 -4.64 -4.86 -4.61 ...
## $ Nativelikeness : num 1.69 2.93 2.48 2.96 1.18 ...
## $ IsIESpeaker   : logi FALSE FALSE FALSE FALSE FALSE ...
## $ LDND         : num 99.6 99.6 99.6 99.6 99.6 99.6 99.6 99.6 99.6 ...
## $ LDN          : num 0.00939 0.00939 0.00939 0.00939 0.00939 ...
## $ PMI          : num 0.827 0.827 0.827 0.827 0.827 ...
## $ Age          : num 32 35 18 23 40 43 19 24 19 43 ...
## $ IsMale        : num 1 1 1 0 1 0 1 1 0 1 ...
## $ AEO          : num 11 15 6 10 7 14 17 9 15 18 ...
## $ LR           : num 8 0 0 0 3 28 0.3 1 0.6 8 ...
## $ NrLang       : int 0 1 1 2 0 0 0 0 1 1 ...
## $ GNI.log      : num 8.06 8.34 10.78 11 8.57 ...
## $ CountryEduYrs : num 5.56 4.37 6.1 9.27 6.4 ...
## $ LDND.c       : num 5.87 5.87 5.87 5.87 5.87 ...
## $ LDN.c        : num 0.000274 0.000274 0.000274 0.000274 0.000274 ...
## $ PMI.c        : num 0.0438 0.0438 0.0438 0.0438 0.0438 ...
## $ Age.c        : num -0.457 2.543 -14.457 -9.457 7.543 ...
## $ AEO.c        : num -1.17 2.83 -6.17 -2.17 -5.17 ...
## $ LR.c         : num 1.18 -6.82 -6.82 -6.82 -3.82 ...
## $ NrLang.c     : num -1.336 -0.336 -0.336 0.664 -1.336 ...
## $ GNI.log.c    : num -0.892 -0.614 1.822 2.046 -0.386 ...
## $ CountryEduYrs.c: num -2.44 -3.64 -1.9 1.27 -1.6 ...

m0n <- gam(Nativelikeness ~ LDND.c + AEO.c + LR.c*Age.c + CountryEduYrs.c + NrLang.c +
            s(Country,bs='re'), data=nat, method='ML')

summary(m0n) # random intercept for country not needed

## 
## Family: gaussian
## Link function: identity
##
## Formula:
## Nativelikeness ~ LDND.c + AEO.c + LR.c * Age.c + CountryEduYrs.c +
##     NrLang.c + s(Country, bs = "re")
##
## Parametric coefficients:
##             Estimate Std. Error t value Pr(>|t|)    
## (Intercept) 2.8268013  0.0642810 43.976 < 2e-16  
## LDND.c      -0.0212581  0.0057234 -3.714 0.000256  
## AEO.c       -0.0452705  0.0096076 -4.712 4.26e-06  
## LR.c        0.0365949  0.0089080  4.108 5.56e-05  
## Age.c       -0.0107250  0.0056574 -1.896 0.059255  
## CountryEduYrs.c 0.0835559  0.0189305  4.414 1.57e-05 
## NrLang.c     0.2839869  0.0456051  6.227 2.25e-09  
## LR.c:Age.c   -0.0011767  0.0004874 -2.414 0.016551  
## 
## Approximate significance of smooth terms:
##             edf Ref.df   F p-value    
## s(Country) 0.2968     80 0.004  0.364
## 
```

```

## R-sq.(adj) = 0.392 Deviance explained = 41.1%
## -ML = 271.63 Scale est. = 0.59885 n = 237

m1n <- gam(Nativelikeness ~ LDND.c + AEO.c + LR.c*Age.c+ CountryEduYrs.c + NrLang.c,
             data=nat, method='ML')

# PMI measure is not better
m1p <- gam(Nativelikeness ~ PMI.c + AEO.c + LR.c*Age.c+ CountryEduYrs.c + NrLang.c,
             data=nat, method='ML')
compareML(m1n,m1p)

## m1n: Nativelikeness ~ LDND.c + AEO.c + LR.c * Age.c + CountryEduYrs.c +
##           NrLang.c
##
## m1p: Nativelikeness ~ PMI.c + AEO.c + LR.c * Age.c + CountryEduYrs.c +
##           NrLang.c
##
## Model m1n preferred: lower ML score (0.748), and equal df (0.000).
## -----
##   Model   Score Edf Difference   Df
## 1   m1p 272.3797   8
## 2   m1n 271.6320   8       0.748 0.000
##
## AIC difference: -1.50, model m1n has lower AIC.

# LDN measure is slightly better
m1l <- gam(Nativelikeness ~ LDN.c + AEO.c + LR.c*Age.c+ CountryEduYrs.c + NrLang.c,
             data=nat, method='ML')
compareML(m1n,m1l)

## m1n: Nativelikeness ~ LDND.c + AEO.c + LR.c * Age.c + CountryEduYrs.c +
##           NrLang.c
##
## m1l: Nativelikeness ~ LDN.c + AEO.c + LR.c * Age.c + CountryEduYrs.c +
##           NrLang.c
##
## Model m1l preferred: lower ML score (0.442), and equal df (0.000).
## -----
##   Model   Score Edf Difference   Df
## 1   m1n 271.6320   8
## 2   m1l 271.1896   8      -0.442 0.000
##
## AIC difference: 0.88, model m1l has lower AIC.

# summary of final model (using LDN)
summary(m1l)

##
## Family: gaussian
## Link function: identity
##

```

```

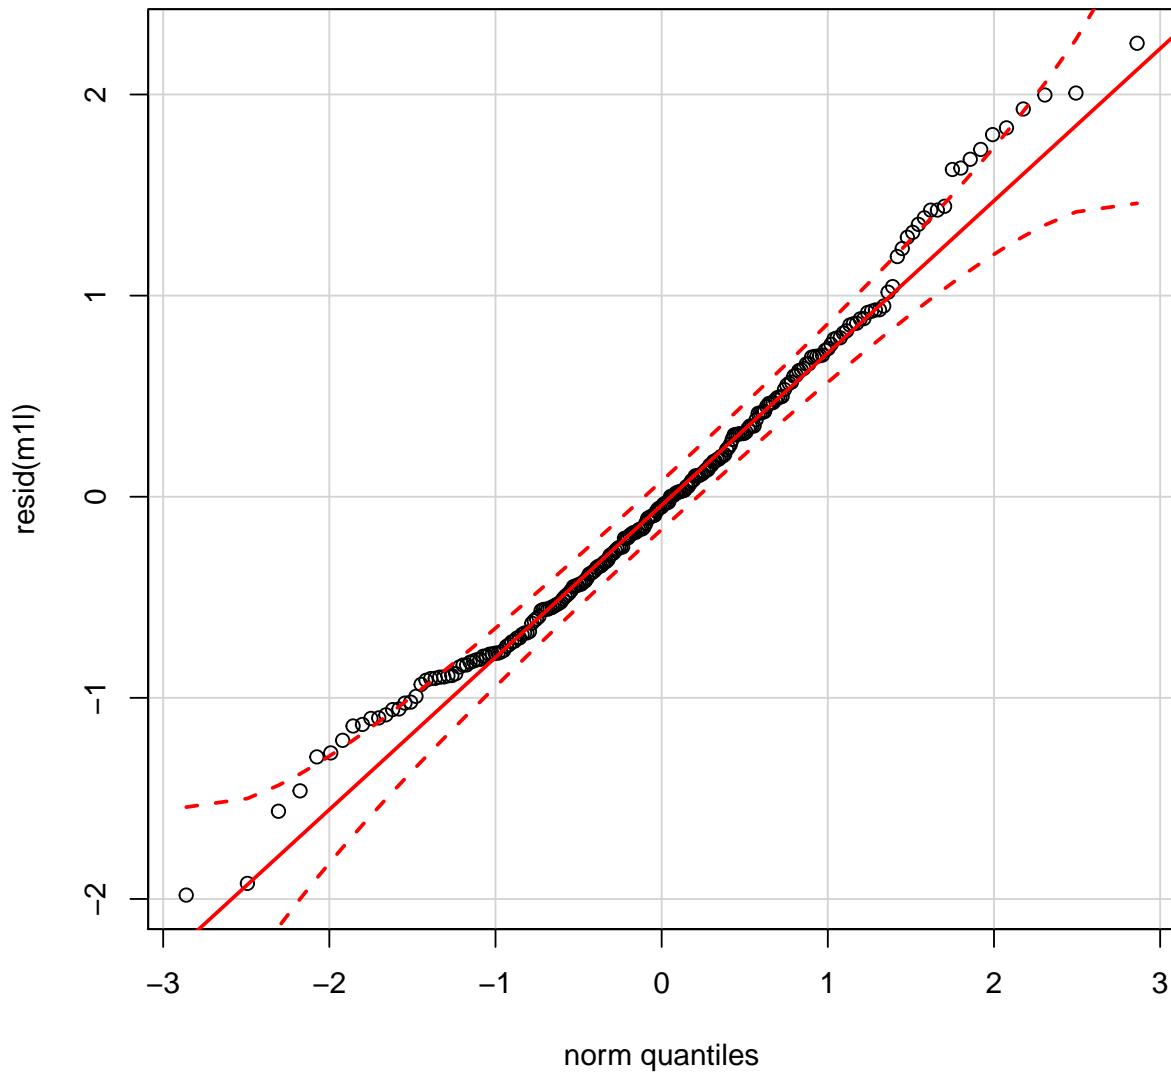
## Formula:
## NativeLikeness ~ LDN.c + AEO.c + LR.c * Age.c + CountryEduYrs.c +
##      NrLang.c
##
## Parametric coefficients:
##                               Estimate Std. Error t value Pr(>|t|)
## (Intercept)            2.832e+00  6.406e-02 44.211 < 2e-16
## LDN.c                  -2.076e+02  5.411e+01 -3.837 0.000161
## AEO.c                  -4.596e-02  9.555e-03 -4.810 2.73e-06
## LR.c                   3.800e-02  8.933e-03  4.254 3.06e-05
## Age.c                 -1.052e-02  5.625e-03 -1.870 0.062820
## CountryEduYrs.c       8.395e-02  1.869e-02  4.491 1.12e-05
## NrLang.c                2.830e-01  4.552e-02  6.217 2.38e-09
## LR.c:Age.c             -1.223e-03  4.865e-04 -2.514 0.012634
##
##
## R-sq.(adj) =  0.393   Deviance explained = 41.1%
## -ML = 271.19   Scale est. = 0.59749   n = 237

# which is similar to the summary of the model
# on the basis of LD.log (but note inverted sign)
summary(m1)

##
## Family: gaussian
## Link function: identity
##
## Formula:
## LD.log ~ LDND.c + AEO.c + LR.c * Age.c + CountryEduYrs.c + NrLang.c +
##      s(Country, bs = "re")
##
## Parametric coefficients:
##                               Estimate Std. Error t value Pr(>|t|)
## (Intercept)            -4.816e+00  8.693e-03 -553.940 < 2e-16
## LDND.c                  3.653e-03  9.134e-04   4.000 7.06e-05
## AEO.c                  9.926e-03  1.020e-03   9.732 < 2e-16
## LR.c                   -5.492e-03  1.034e-03  -5.313 1.47e-07
## Age.c                  2.143e-03  6.808e-04   3.147 0.001723
## CountryEduYrs.c      -1.971e-02  3.013e-03  -6.539 1.23e-10
## NrLang.c                -1.936e-02  5.763e-03  -3.359 0.000826
## LR.c:Age.c              1.077e-04  4.025e-05   2.675 0.007659
##
## Approximate significance of smooth terms:
##          edf Ref.df    F p-value
## s(Country) 35.67     130 0.505 1.55e-05
##
## R-sq.(adj) =  0.368   Deviance explained = 40.6%
## -ML = -277.77   Scale est. = 0.025279   n = 711

# residuals are approximately normally distributed
qqp(resid(m1))

```



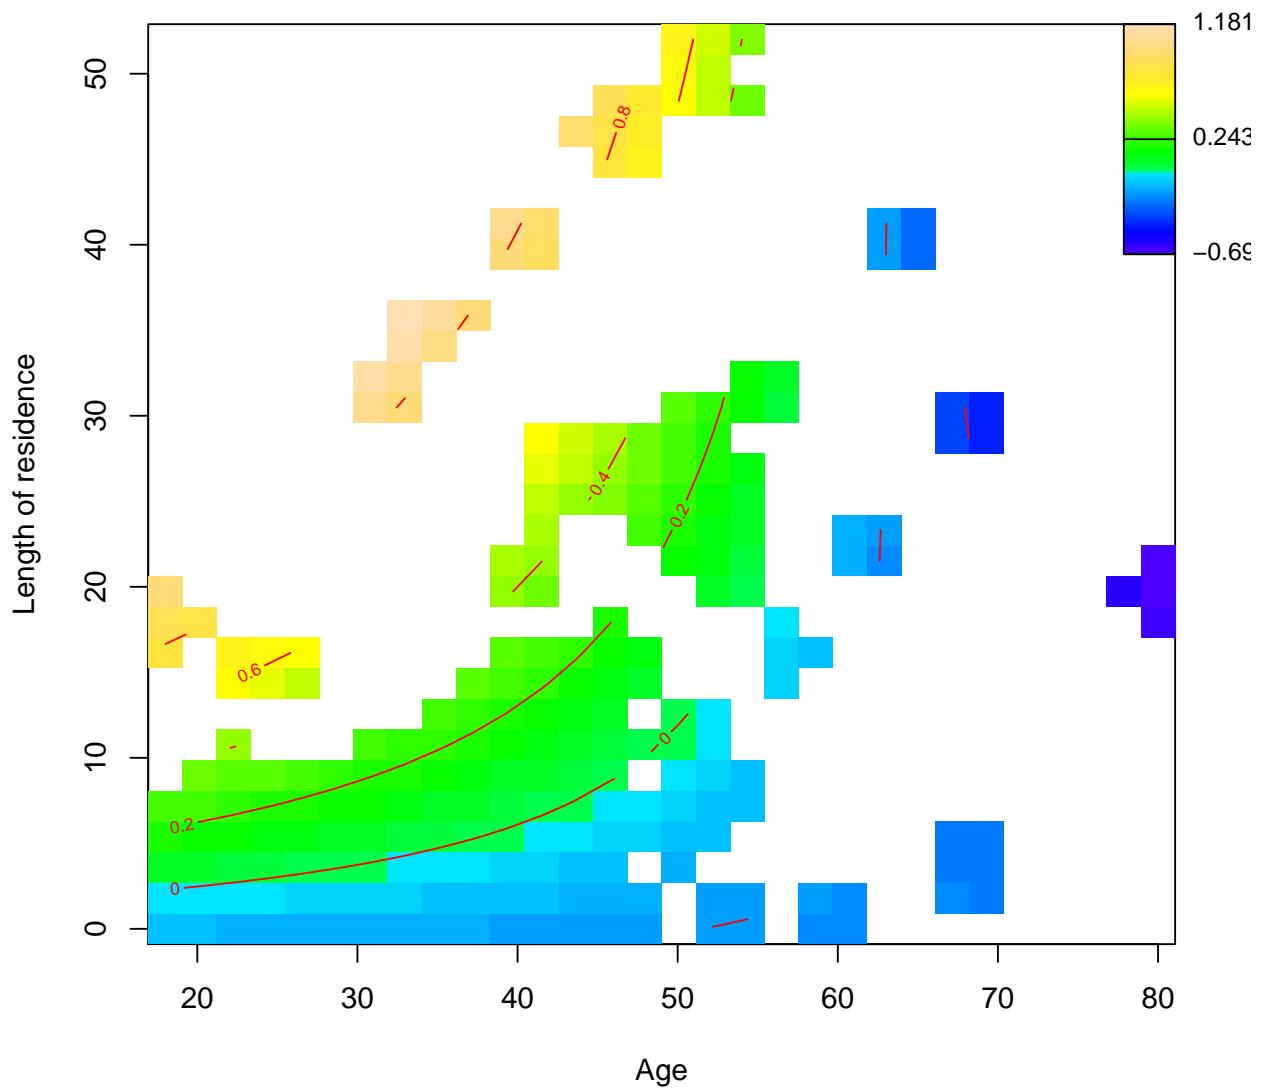
```

# visualization of interaction (refitting with tensor,
# which in this case simply models a linear interaction
# as the number of DF equals 3)
m1lb <- gam(Nativelikeness ~ LDN + AEO + te(Age,LR) + CountryEduYrs + NrLang +
  s(Country,bs='re'), data=nat, method='ML')

pvisgam(m1lb, select=1, view=c('Age','LR'), too.far=0.04, hide.label=T,
  ylab='Length of residence', xlab='Age', main='')

## [1] "Tensor(s) to be plotted: te(Age,LR)"

```



```
# explained variance of LDN measure by itself
summary(gam(Nativelikeness ~ LDN.c, data=nat, method='ML'))

##
## Family: gaussian
## Link function: identity
##
## Formula:
## Nativelikeness ~ LDN.c
##
## Parametric coefficients:
```

```

##             Estimate Std. Error t value Pr(>|t|)
## (Intercept)  2.77146   0.06034 45.933 < 2e-16
## LDN.c      -332.85477  56.59837 -5.881 1.39e-08
##
## 
## R-sq.(adj) =  0.125   Deviance explained = 12.8%
## -ML = 317.72   Scale est. = 0.86221   n = 237

# despite the non-significance of age, the interaction is still necessary
m1lc <- gam(Nativelikeness ~ LDN.c + AEO.c + LR.c + Age.c + CountryEduYrs.c + NrLang.c,
  data=nat, method='ML')
compareML(m1l,m1lc) # m1l is better

## m1l: Nativelikeness ~ LDN.c + AEO.c + LR.c * Age.c + CountryEduYrs.c +
##           NrLang.c
##
## m1lc: Nativelikeness ~ LDN.c + AEO.c + LR.c + Age.c + CountryEduYrs.c +
##           NrLang.c
##
## Chi-square test of ML scores
## -----
##    Model     Score Edf Chisq    Df p.value Sig.
## 1  m1lc 274.4150    7
## 2  m1l 271.1896    8 3.225 1.000   0.011  *
##
## AIC difference: -4.45, model m1l has lower AIC.

```