

Strategies for addressing collinearity in multivariate linguistic data

Fabian Tomaschek, Peter Hendrix, and R. Harald Baayen

Department of General Linguistics, University of Tübingen, Germany

July 18, 2018

Abstract

When multiple correlated predictors are considered jointly in regression modeling, estimated coefficients may assume counterintuitive and theoretically uninterpretable values. We survey several statistical methods that implement strategies for the analysis of collinear data: regression with regularization (the elastic net), supervised component generalized linear regression, and random forests. Methods are illustrated for a data set with a wide range of predictors for segment duration in a German speech corpus. Results broadly converge, but each method has its own strengths and weaknesses. Jointly, they provide the analyst with somewhat different but complementary perspectives on the structure of collinear data.

keywords: elastic net, supervised component generalized linear regression, random forests, collinearity, concurvity, segment duration.

25 1. Introduction

26 Response measures in linguistics and phonetics are often a function not of a single pre-
27 dictor but of many predictors jointly, reflecting a move away from mono-causal to multi-
28 factorial explanations. For instance, reductions and deletions in speech have been shown
29 to correlate with a range of measures which include frequencies of occurrence and condi-
30 tional probabilities at word and segment level (among others Jurafsky et al., 2000; Aylett
31 and Turk, 2004; Gahl, 2008; Bell et al., 2009; Tremblay and Tucker, 2011; Priva, 2015).
32 For example, Tremblay and Tucker (2011) used no less than 18 such measures to predict
33 the durations of four-word sequences. Typically, many of the covariates included in these
34 analyses serve as controls for potential confounds with predictors of central theoretical
35 interest.

36 When predictors are completely uncorrelated and fully orthogonal, the results of a
37 multivariable regression model and separate regressions with one predictor each will be
38 virtually identical. Multiple regression comes into its own for data with non-orthogonal
39 predictors. For such data, it serves as a mathematically principled arbiter for teasing
40 apart relevant from irrelevant predictors. However, when predictors are strongly corre-
41 lated, i.e., for collinear data, this arbitrage tends to result in counterintuitive and unin-
42 terpretable coefficients (Farrar and Glauber, 1967; Belsley et al., 1980). In this study, we
43 review statistical methods that work around this problem.

44 When a data set is characterized by substantial collinearity, several problems arise.
45 First, as already mentioned, parameter estimates may assume unexpected and theoret-
46 ically uninterpretable values. Second, the model fit to the data will be unstable, in the
47 sense that removal of just a few data points may have substantial consequences for the
48 estimates of regression parameters. This holds both for linear regression and for the linear
49 mixed model. Third, it can happen that no predictor on its own is significant, whereas
50 all predictors jointly are successful in explaining a significant part of the variance in the
51 response (Chatterjee et al., 2000).

52 In what follows, we begin with an introduction to the problem of collinearity¹ and

¹In the context of nonlinear regression, collinearity also rears its ugly head in the form of concurvity. Concurvity can render models such as generalized additive (mixed) models unstable. We therefore briefly discuss how concurvity can be assessed, and what measures the analyst might consider when concurvity is high, in the appendix.

53 its adverse consequences for the magnitude and sign of estimated coefficients. We then
54 describe a data set with substantial collinearity that will serve as the test case for our
55 analyses. Subsequently, we introduce and illustrate three methods for analyzing collinear
56 data. The first of these is a non-parametric technique from machine learning, random
57 forests. Random forests enable the analyst to assess the relative importance of predic-
58 tors. The second method is supervised component generalized linear regression (SCGLR).
59 SCGLR performs dimensionality reduction on the predictor space, resulting in a smaller
60 set of orthogonal predictors (the supervised components). SCGLR comes with visualiza-
61 tion methods for inspecting how the original predictors load on the supervised compo-
62 nents, and it provides regression coefficients for the original predictors that are properly
63 shrunk. The third method that we discuss is the elastic net, a regularized regression
64 technique that not only shrinks coefficients, but shrinks some of these completely to zero.
65 This method therefore can be used to perform variable selection. For each method, we
66 introduce the general concepts, and then illustrate its use for our example data set.

67 There is no fixed set of guidelines that guarantee the “correct” analysis of collinear
68 data. George Box’s famous aphorism that all models are wrong but some are useful (Box,
69 1976) is especially relevant with respect to models for highly collinear data. The methods
70 we review in the present study therefore provide the analyst with a toolkit that we find
71 useful for exploring and understanding in complementary ways to what extent, and how
72 a response might be shaped by a set of collinear predictors.

73 All analyses discussed in this study are documented step by step in the supplementary
74 materials, to be downloaded from <https://osf.io/5merb/>. For these analyses, we made
75 use of the statistical programming environment R (R Core Team, 2018) and specialist
76 packages available for R (introduced below).

77 **2. Suppression and enhancement**

78 Suppression and enhancement occur in the linear regression model when two (or more)
79 predictors for a given response Y are strongly correlated. Take, for example, an analysis
80 in which response times (dependent variable Y) in auditory lexical decision have to be
81 predicted by word frequency counts in American English (predictor A) and British English
82 (predictor B). Given that such frequency counts will tend to be strongly correlated,
83 suppression and enhancement are likely to make the coefficients of the regression model

84 uninterpretable. To understand why this happens, first consider the case in which we fit
 85 two one-predictor regression models to Y ,

$$Y_i = \beta_0 + \beta_A A_i + \epsilon_i, \epsilon_i \sim \mathcal{N}(0, \sigma), \quad (1)$$

$$Y_i = \beta_0 + \beta_B B_i + \epsilon_i, \epsilon_i \sim \mathcal{N}(0, \sigma). \quad (2)$$

where the β_0 represent the intercepts, β_A and β_B denote the coefficients for predictors A and B , and ϵ is a Gaussian error term. When A and B are uncorrelated and completely orthogonal, the results of these two one-predictor models will almost completely identical to a multivariable regression model in Y in predicted from A and B jointly:

$$Y_i = \beta_0 + \beta_A A_i + \beta_B B_i + \epsilon_i, \epsilon_i \sim \mathcal{N}(0, \sigma). \quad (3)$$

86 In this case, the multivariable regression model has nothing to add about the effects of
 87 A and B that we did not already know from the two one-predictor analysis. However,
 88 when A and B are correlated, and not strictly orthogonal, then multiple regression comes
 89 into its own as the arbiter deciding which predictors should be given more (or less)
 90 weight. When predictors are only mildly correlated, there is little collinearity and the
 91 weights estimated by the multiple regression model (3) will make sense, but when strong
 92 collinearity is present, the resulting model will become theoretically uninterpretable.

93 Following Friedman and Wall (2005), we illustrate this phenomenon by varying the
 94 correlation between predictors A and B , while keeping constant the correlations between
 95 A and Y as well as the correlations between B and Y . We simulated multiple data
 96 sets with 1000 observations each, using the `mvrnorm` function from the **MASS** package
 97 (Venables and Ripley, 2002). Y , A and B are all standard normal random variables. We
 98 manipulated the correlation between A and B (r_{AB}) to range from close -1 to close to $+1$
 99 in steps of 0.01. We fixed the correlation between B and Y at $r_{BY} = 0.3$, but considered
 100 three different correlations between A and Y : $r_{AY} = -0.3$, $r_{AY} = 0.0$ and $r_{AY} = 0.6$.
 101 When $r_{AB} = 0$, β_A is equal to r_{AY} and $\beta_B = r_{BY}$.

102 Figure 1 illustrates the consequences of varying the correlation between A and B for
 103 the estimates of slopes β_A and β_B (top panels) and the corresponding t -values (bottom
 104 panels). Across all panels of Figure 1, dashed lines represent β_A and solid lines β_B . The
 105 three values of r_{AY} are listed above their respective panels.

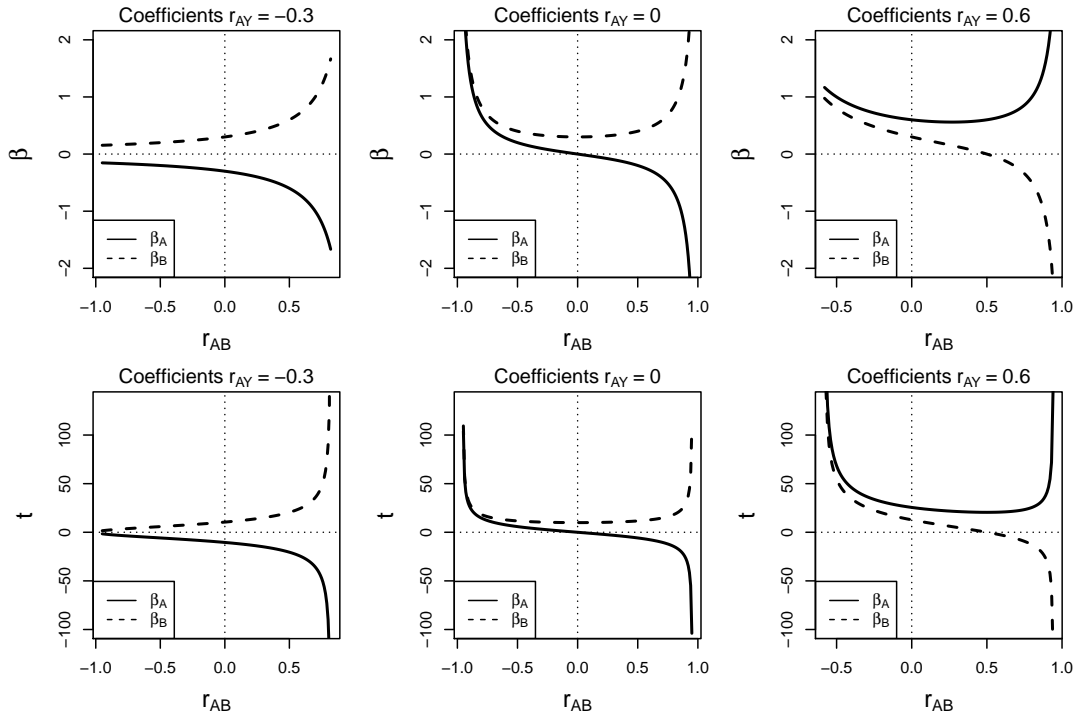


Figure 1: β coefficients (top row), and t -values (bottom row) as a function of r_{AB} , for $r_{BY} = 0.3$ and varying correlations r_{AY} (left column: -0.3 , middle column: 0 , right column: 0.6).

106 First consider the panels graphing coefficients against r_{AB} . When r_{AB} is zero, β_A is
 107 -0.3 when $r_{AY} = -0.3$, it is 0 when $r_{AY} = 0$, and it is 0.6 when $r_{AY} = 0.6$. As r_{BY}
 108 is fixed at 0.3 , β_B is always 0.3 when $r_{AB} = 0$. When r_{AB} moves away from zero, the
 109 coefficients change, and the more extreme r_{AB} becomes, the more extreme the changes
 110 in the coefficients are. When r_{AB} approximates 1 , we find large positive and negative
 111 values for both β_A and β_B . Which predictor receives a positive coefficient and which a
 112 negative depends on r_{AB} . When r_{AB} is shifted towards -1 , coefficients are not enhanced,
 113 but suppressed: both β_A and β_B assume smaller values than they have when $r_{AB} = 0$. It
 114 is noteworthy that β_A is strongly enhanced even when $r_{AY} = 0$.

115 Estimates of the t -values associated with the coefficients also vary with r_{AB} and can
 116 be very large for extreme positive values of r_{AB} . This leads to false positives for β_A
 117 when $r_{AY} = 0$ and r_{AB} is large. In other words, the model supports a significant effect
 118 of A although there is in fact none. False negatives arise when $r_{AY} = -0.3$, $r_{BY} = 0.3$,
 119 and r_{AB} is close to -1 . In other words, the model does not support a significant effect
 120 of A and B although they are in fact significantly correlated with Y . In fact, strong
 121 collinearity can give rise to a model that succeeds in explaining variance of the predictor,

122 without a single regressor being significant (see, e.g., Hadi, 1988; Chatterjee and Hadi,
123 2012b; Friedman and Wall, 2005, for examples).

124 Large coefficients with opposite sign for strongly correlated predictors are the hallmark
125 of collinearity. In this case, the coefficients become difficult to interpret. For the above
126 example of American and British frequency of occurrence, one frequency measure will
127 reveal a coefficient with the expected negative sign, but the other frequency measure will
128 emerge with a coefficient with an uninterpretable positive sign.

129 When strong collinearity is present, it is important to take a step back, and to address
130 the question of how the artifacts of strong collinearity are best avoided. Before introducing
131 possible strategies for addressing the adverse effects of collinearity, we first introduce the
132 data set that we use to illustrate these strategies, the KIEL corpus.

133 **3. Data set: word and segment durations in the KIEL corpus**

134 The KIEL corpus (Kohler, 1996; Peters, 2003) comprises quasi-spontaneous speech as well
135 as speech elicited by dictation. The corpus is annotated at the word level, the segment
136 level, and the prosodic level. Annotations at the segmental level were manually corrected
137 and contain indicators about missing canonical segments. Prosodic annotation provides
138 information about primary and secondary stress in words. The entire corpus contains
139 32,460 word tokens (2,216 types), recorded from a total of 107 speakers.

140 From the KIEL corpus, we extracted durations for those vowels that occur in mono-
141 syllabic words and that were recorded in quasi-spontaneous speech. Of this set of vowels,
142 we selected the first 10,000 (from a total of 314 unique word types) for further analysis.
143 The response variable of interest is vowel duration.

144 For each vowel, we registered `speaker`, `carrier word`, and `segment` identity, three
145 random-effect factors. We recorded `stress` (levels `none`, `primary`, `secondary`), an
146 indicator variable for whether the segment is located in a word at the end of a sen-
147 tence, (`EndOfSentence`, with levels `true`, `false`), and phonological length of the vowel
148 (`VowelLength`, with levels `long`, `short`).

149 In addition, we included `SpeakingRate` (number of syllables per second) and word du-
150 ration (`wDur`). Following previous research (Jurafsky et al., 2000; Aylett and Turk, 2004;
151 Bell et al., 2009; Tremblay and Tucker, 2011; Priva, 2015), we added 16 probabilities for
152 segments and words from the frequencies of words and segments in the KIEL corpus. In

153 what follows, we use **W** to denote words, **S** to denote segments, **target** for the current unit
154 (**W** or **S**), and **prev** and **next** to denote preceding and following units. The probabilities
155 we considered are: the probability (relative frequency) of the preceding, current, and fol-
156 lowing unit: $P(W_{\text{target}})$, $P(W_{\text{next}})$, $P(W_{\text{prev}})$, $P(S_{\text{target}})$, $P(S_{\text{prev}})$, $P(S_{\text{next}})$; the
157 joint probability with the preceding, or following unit: $P(W_{\text{prev}}, W_{\text{target}})$, $P(W_{\text{target}},$
158 $W_{\text{next}})$, $P(S_{\text{prev}}, S_{\text{target}})$, $P(S_{\text{target}}, S_{\text{next}})$; the joint probability with both the
159 preceding and following unit: $P(W_{\text{prev}}, W_{\text{target}}, W_{\text{next}})$, $P(W_{\text{prev}}, W_{\text{target}}, W_{\text{next}})$;
160 the conditional probability given the preceding unit: $P(W_{\text{target}} | W_{\text{prev}})$, $P(S_{\text{target}} |$
161 $S_{\text{prev}})$; and the conditional probabilities given the following unit: $P(W_{\text{target}} | W_{\text{next}})$,
162 $P(S_{\text{target}} | S_{\text{next}})$.

163 To this set of continuous predictors we added a final set of covariates: phonological
164 neighborhood density (**NHD**), the count of words identical to the target word except for one
165 segment; the count of segments in a word (**nSegperWord**); the number of speakers using a
166 word (**Dispersion**) (see Adelman et al., 2006; Keuleers et al., 2015, for lexical dispersion
167 across texts and speakers). In recent years, more and more researchers use measures
168 derived from cognitive and neural networks to predict human behavior in cognitive tasks.
169 These measures, such as activation estimated with naive discriminative learning (Baayen
170 et al., 2011, 2016; Milin et al., 2017), are correlated with frequency measures to vari-
171 ous degrees. To increase the number of potentially correlated predictors, we added the
172 activation of the word given a word’s diphones (**WordActivation SmallWindow**), and
173 the activation of the word as provided by all diphones that occur in a five-word win-
174 dow around the target word (see Tomaschek et al., 2018, for further discussion). Larger
175 activations are expected to be associated with shorter durations. The total number of
176 numeric predictors thus amounts to 24. There are potentially other collinear predictors
177 due to the nature of how they were created. For example, conditional probabilities are
178 derived from frequencies of occurrence, which is why they are collinear by design.

179 Before analysis, we transformed numeric variables where necessary. As indicated by
180 a Box-Cox test, the response variable was transformed by taking its square root. Several
181 predictors were subjected to either a logarithmic transform, or to the root transform,
182 depending on which transformation succeeded in rendering the distribution of values more
183 symmetrical and with fewer outliers. For discussion of why transformations of response
184 and predictor variables are necessary in the context of linear regression, see Zuur et al.




Figure 2: Correlation map for numeric predictors in the KIEL corpus.

185 (2010); Chatterjee and Hadi (2012a); Sheather (2009).

186 4. Diagnostics for collinearity

187 4.1. Correlation plot

188 When a linear model is fit to the segment durations, a first indication of trouble is
 189 that there are predictors for which the coefficients are not estimated. Furthermore, which
 190 predictors are inestimable depends on the order of the predictors in the model formula.

191 As a first step towards a diagnosis of what is wrong, we inspect the correlations be-
 192 tween the predictors, using a correlation map (obtained with the **corrplot** package (Wei
 193 et al., 2017)). In Figure 2, red dots represent positive correlations, whereas blue dots
 194 represent negative correlations. The size of the dots is proportional to the magnitude of
 195 the correlation. It is clear that many predictors are correlated to some extent. There are
 196 especially large correlations for **WordActivation SmallWindow** and **P(Sprev, Starget,**
 197 **Snext)**: $r = 0.77$, **Dispersion** and **P(Wtarget)**: $r = 0.90$, **nSegperWord** and **NHD**: $r =$

198 -0.78, `WordActivation SmallWindow` and `WordActivation LargeWindow`: $r = 0.91$, and
199 $P(\text{Starget} \mid \text{Snext})$ and $P(\text{Starget}, \text{Snext})$: $r = 0.76$. The problem with the correla-
200 tion matrix as a diagnostic for collinearity is that although high correlations indeed point
201 to a potential collinearity problem, the absence of high correlations does not guarantee
202 that there is no problem (see Belsley et al., 1980, p. 92–93 for further discussion).

203 4.2. Variance inflation factors

A better diagnostic for assessing whether coefficients are poorly estimated due to
collinearity are the variance inflation factors (VIF) for the coefficients. The variance
 $\text{VAR}[\hat{\beta}_j]$ of an estimated coefficient $\hat{\beta}_j$ for predictor j is

$$\text{VAR}[\hat{\beta}_j] = \frac{1}{1 - R_j^2} \cdot \frac{\sigma^2}{(n - 1)S_j^2}, \quad (4)$$

204 where S_j denotes the standard deviation of predictor j , n is the number of data points, σ^2
205 the common variance of the errors, and R_j^2 the value of R^2 obtained from regressing the
206 j -th predictor on all other remaining predictors. When predictor j is highly dependent
207 on one or more other predictors, R_j^2 will be large, and as a consequence $1/(1 - R_j^2)$ will
208 be large as well. If predictor j is orthogonal to the other predictors, R_j^2 is close to zero,
209 and $1/(1 - R_j^2)$ close to 1. The ratio $1/(1 - R_j^2)$ is called the j -th variance inflation factor.
210 One rule of thumb is that coefficients with a variance inflation factor exceeding five are
211 poorly estimated and untrustworthy (Sheather, 2009). In R, variance inflation factors
212 can be obtained with, e.g., the `vif()` function of the `car` package (Fox and Weisberg,
213 2011). When we try to apply `vif()` to the above-mentioned linear model, it reports
214 that it cannot do so: not all coefficients in this model are estimable. When we refit the
215 model with two troublesome predictors removed (e.g., `P(Wnext)`, `P(Wprev)`), we find that
216 there are 13 predictors with variance inflation factors exceeding 5. For five of these, the
217 variance inflation factor exceeds 10.

218 A problem with variance inflation factors is that it is not clear what a meaningful
219 boundary is for a low versus a high value. For instance, Chatterjee and Hadi (2012b)
220 state that values exceeding 10 are diagnostic of collinearity problems (p. 250), whereas
221 (Sheather, 2009) puts the boundary at 5 (p. 203). For the present data, however, it is
222 clear that there is a serious collinearity problem.

Whereas variance inflation factors are useful for finding individual predictors that clearly suffer from collinearity, the collinearity of the full set of predictors jointly is still not well assessed. This led Belsley et al. (1980) to propose a ‘systemic’ measure for collinearity, called the condition number κ . To understand what κ actually assesses, we write out the estimates of the coefficients as a function of the model matrix \mathbf{X} (the matrix with the predictors and a column of ones for the intercept) and the observed values of the response \mathbf{y} :

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

The Achilles heel of the linear model is calculating the inverse of the (square) matrix $\mathbf{X}^T \mathbf{X}$. (When all predictors are centered and scaled, $\mathbf{X}^T \mathbf{X}$ is the correlation matrix.) The inverse of a square matrix need not exist. It does not exist when there are columns (or rows) that are (weighted) combinations of each other. To ascertain whether a matrix is singular, it can be decomposed into a product of three matrices, the middle matrix of which is zero except possibly for the elements on its main diagonal. These elements are known as eigenvalues. When a matrix is singular, at least one of these eigenvalues is zero. For empirical data, it is unlikely that predictors will be exactly (weighted) combinations of each other. Typically, predictors are themselves not exact but noisy.

Nevertheless, the more similar one or more empirical predictors are, the more $\mathbf{X}^T \mathbf{X}$ starts to resemble a singular matrix. This resemblance becomes stronger when one or more eigenvalues of $\mathbf{X}^T \mathbf{X}$ are very close to zero. As we need the reciprocals of the eigenvalues to calculate the inverse matrix, it is clear that eigenvalues close to zero are going to give rise to huge reciprocals. Such huge reciprocals make the inverse matrix, and hence the estimates of the coefficients, unstable.

It turns out that the eigenvalues of $\mathbf{X}^T \mathbf{X}$ are the squares of the so-called singular values of the design matrix \mathbf{X} (the diagonal elements of the center matrix when \mathbf{X} itself is decomposed into a product of three matrices). Therefore, very small singular values for \mathbf{X} are also indicative of a collinearity problem. Belsley et al. (1980) show that the ratio of the largest and smallest singular values, the condition number κ , is the pivotal scaling factor for an upper bound for the effect of small changes in the response variable on the magnitude of the coefficients. Likewise, it provides such a scaling factor for small changes

246 in the predictors. In other words, if κ is large, very small differences in the response or
247 predictor variables have huge consequences for the estimated regression coefficients.

248 For calculating κ , we start off with the matrix with predictors, we add a column of
249 ones for the intercept, and then scale each column so that it has unit length. Without this
250 scaling, the value of κ would depend heavily on the measurement units of the variables,
251 and as a consequence, it would become useless as a general diagnostic of collinearity.
252 Belsley et al. (1980) point out that predictors should not be centered (see also Belsley,
253 1984, for detailed discussion): transformation of variables to Z-scores does not remove
254 collinearity but makes it invisible. The singular values of the resulting matrix can be
255 calculated, from which we obtain κ . All this is implemented in `collin.fnc()` from the
256 **languageR** package (Baayen, 2008), which follows Belsley et al. (1980). (The `kappa`
257 function of R does not include the intercept, and hence, even when its directive `exact`
258 is set to `TRUE`, will give rise to different results.) Values of κ exceeding 15 typically
259 indicate that harmful effects of collinearity will be present. Values exceeding 30 point to
260 strong collinearity for which corrective action is essential. These cutoff values are based
261 on experience that has accumulated over the years in data analysis (Belsley et al., 1980;
262 Chatterjee and Hadi, 2012b). For the predictors in the KIEL data set, κ is no less than
263 1, 809, 457, 843, 187, 094.

264 4.4. *Inspecting the sign*

265 When in doubt about the severity of collinearity and potential adverse effects of
266 enhancement, it may be useful to check whether the sign of a coefficient is in accordance
267 with the sign of a simple correlation of the same predictor with the response. If there is
268 indeed a change of sign, it is worth investigating whether corrective measures are required.

269 5. Strategies for addressing collinearity

270 5.1. *Common sense strategies*

271 When the set of predictors includes a set of variables that are theoretically strongly
272 related, it makes sense to include only one in the regression analysis. By way of example,
273 frequency counts based on a range of corpora will show strong correlations. When the
274 nature of these corpora and the corresponding consequences for word use are not of

275 primary interest, selecting one frequency measure from the set will help bring down
276 collinearity.

277 Instead of selecting one measure by hand, one could alternatively orthogonalize the
278 available measures using, for instance, principal component analysis, and select the first
279 principal component, or the first couple of principal components, as representative for
280 the full set of measures. Principal component analysis is explained in more detail in
281 Section 5.3.1. Baayen et al. (2006) used this approach for 10 strongly correlated measures
282 of orthographic and phonological consistency. Below, we discuss a method, supervised
283 component generalized linear regression, that carries out orthogonalization in a more
284 principled way.

285 Sometimes it is possible to de-correlate two related predictors by selecting one pre-
286 dictor and including the ratio of the first and second predictor as a new predictor. For
287 instance, Baayen et al. (2006) were interested in frequency of occurrence in spoken and
288 written English, and included written English as one predictor, and the ratio of written
289 to spoken English as second predictor. The new predictor, which gauges the extent to
290 which a word is used more often in writing than in speech, is by far not as strongly
291 correlated with written frequency as the original spoken frequency measure.

292 These common-sense strategies all share one disadvantage: a strong dependence on
293 manual intervention. Although hand-crafting the set of predictors may be justified by
294 domain knowledge, methods that minimize manual intervention are worth considering.
295 We discuss three such methods below.

296 One strategy that is not recommended is to reduce collinearity through residualization
297 (see, e.g. Tremblay and Tucker, 2011; Priva, 2015, for applications of this strategy). A
298 predictor A that is correlated with another predictor B is not entered into the analysis
299 directly. Instead, A is regressed against predictor B , and the residuals of this regression
300 ($A_{\text{residuals}}$) are then entered into the analysis as a predictor instead of A . Since $A_{\text{residuals}}$
301 is orthogonal to predictor B , this reduces collinearity.

302 York (2012) and Wurm and Fisicaro (2014), however, demonstrated that the statistical
303 characteristics of β_A and $\beta_{A_{\text{residuals}}}$ are identical. By contrast, unfortunately, residualiza-
304 tion can lead to an exaggeration of the statistical importance of the non-residualized
305 predictor B or an overestimation of the importance of data in regions of enhancement,
306 depending on magnitude and sign of the correlation between A and B . As a consequence,

307 residualization may strongly affect the results and the interpretation of a regression anal-
308 ysis. In what follows, we consider strategies for analysing collinear data that do not
309 require removing or orthogonalizing predictors by hand.

310 *5.2. Random forests*

311 Collinearity is a problem of the linear model and the way in which it estimates regres-
312 sion coefficients. One way in which one can avoid the problems that arise in the context
313 of the linear model due to collinearity is to step away from the regression framework,
314 and to use instead a non-parametric method from machine learning. In what follows, we
315 discuss random forests, which make use of decision trees and recursive partitioning.

316 Conditional variable importance measures calculated in random forests take into ac-
317 count the correlations between predictors. One issue with conditional variable impor-
318 tances, however, is that they are heavy on resources. Furthermore, these measures tend
319 to inflate variable importance scores for uncorrelated data (Nicodemus et al., 2010). For
320 this reason, we decided to use the unconditional variable importances provided by the
321 **ranger** package (Wright and Ziegler, 2017) for R^2 .

322 Before discussing further details, we clarify the contexts in which this method is of use.
323 When the aim of the analysis is a model with outstanding prediction accuracy, random
324 forests are an excellent choice. Random forests, however, do not provide detailed insight
325 in the effects of individual predictors and their interactions. What they do provide is
326 an assessment of predictor importance. When interest resides primarily in the effects of
327 individual predictors and their significances, random forests remain useful as a tool for
328 exploratory data analysis, just like visualization.

329 *5.2.1. Recursive partitioning*

330 Random forests are based on decision trees, which use a set of binary rules to predict a
331 response variable. The response variable in a decision tree can be categorical or numerical
332 in nature. Recursive partitioning trees for categorical responses are known as classification
333 trees, trees for numerical responses are referred to as regression trees. The dependent

²We are thankful to Bodo Winter to pointing us to the **ranger** package, which outperforms alternative R packages such as **party** (Hothorn et al., 2018a), **partykit** (Hothorn et al., 2018b), and **randomForest** (Breiman et al., 2018) in terms of computational efficiency.

334 variable in the current study is segment duration, which is numerical, and consequently
335 the decision trees introduced here are regression trees. In the analyses that follow, we
336 use the variables as transformed for regression modeling, but such transformations are
337 not required for analyses based on random forests and decision trees.

338 Decision trees are built through a process that is commonly referred to as recursive
339 partitioning. Recursive partitioning algorithms start with the full data set, which includes
340 all observations. The algorithm starts off with finding the predictor and the predictor
341 value that split the data into two groups in an optimal manner. A commonly used
342 splitting criterion, also used in the random forest analyses below, is the reduction in
343 uncertainty (i.e., the reduction in entropy, which is also referred to as the information
344 gain) about the value of the response variable (e.g., Therneau et al., 2017). Splits are
345 implemented for a predictor value that reduce the uncertainty about the response variable
346 the most. For each of the two subsets of the data that result from the split this process
347 is repeated. The process of implementing binary splits for a branch of the tree continues
348 until a stopping criterion is reached that is based on the extent to which additional splits
349 improve the quality of the model fit. The model fitting procedure is concluded when the
350 stopping criterion has been reached for all branches of the decision tree.

351 An example of a recursive partitioning tree is shown in Figure 3, top. For ease of
352 illustration, we limited this tree to a maximum depth (i.e., number of splits) of 2. The
353 initial split is made on word duration (`wDur`), at a value of -0.51 . Observations for which
354 the (normalized) word duration is smaller than -0.51 are assigned to the left branch
355 of the trees, whereas observations for which word duration is equal to or greater than
356 -0.51 are assigned to the right branch of the tree. The second split in the left branch of
357 the tree is based on the value of phonological neighborhood density (`NHD`), whereas the
358 second split in the right branch of the tree is based on the number of segments in a word
359 (`nSegPerWord`).

360 The colored boxes provide more information about the observations in a node. The
361 top value in a colored box is the mean segment durations for the observations in the
362 corresponding node, whereas the bottom value in a box provides the percentage of obser-
363 vations in the data set that fall under the corresponding node. Mean segment durations
364 for the observations in the four terminal nodes, i.e. the nodes at the last layer of the tree,
365 differ substantially, which demonstrates that the implemented splits were successful at

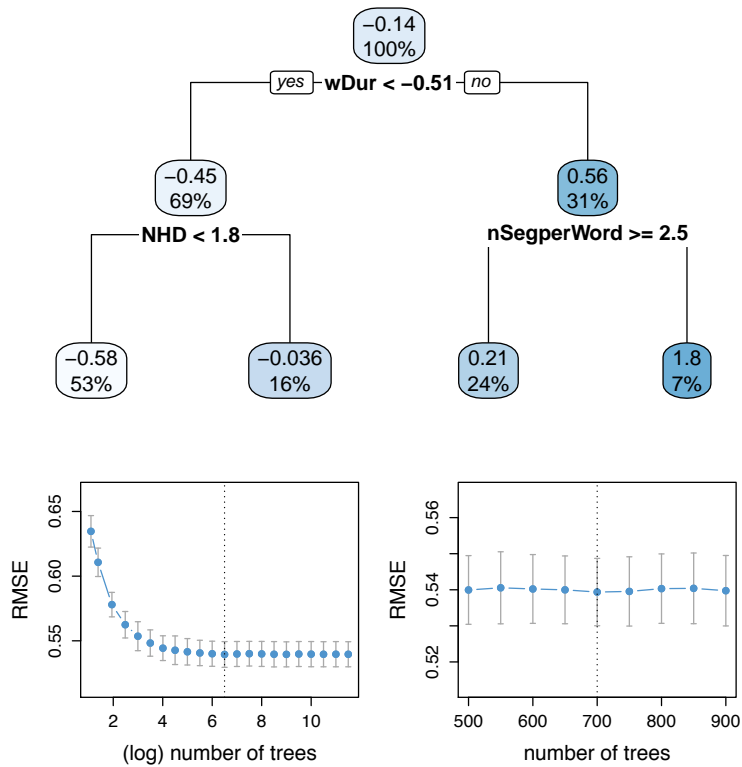


Figure 3: **top:** Recursive partitioning tree fit to the segment durations in the KIEL corpus data. Colored boxes indicate mean predictor values and percentage of observations for the observations in each node. **bottom:** Results of the random forest models fit to the segment durations in the KIEL data. Optimal cross-validation performance for different numbers of trees on a coarse grid (left panel) and on a fine grid (right panel). The dashed lines indicates the number of trees for which the MSE is minimal.

366 dividing the data into subsets with different segment durations.

367 Random forests (Breiman, 2001) fit not one, but multiple decision trees to the data.
 368 The idea behind random forests is to prevent overfitting by averaging over the predictions
 369 of a large number of trees. To make this idea work, it is crucial to ensure that the individ-
 370 ual trees are not too similar. Simply fitting multiple decision trees to the complete data
 371 set would result in a series of identical trees. To overcome this problem, random forests
 372 combine two statistical concepts: bootstrap aggregating (bagging) and random predictor
 373 subset selection. Both of these techniques reduce the correlation between individual trees.

374 Bootstrap aggregating (bagging) is a method to artificially obtain more samples than
 375 the data can provide. The original data set acts as a pseudo-population. From this

376 population, we take pseudo-samples that have the same size as the population and that
377 are drawn from the population with replacement. As a result, a sample contains ap-
378 proximately two thirds of the observations in the population, whereas one third of the
379 observations is left out. The observations that are in the sample are referred to as the
380 in-bag observations, whereas the observations that are not in the sample are referred to as
381 the out-of-bag observations. Each tree in a random forest is fit to a different bootstrapped
382 sample.

383 The trees in random forest are not only fit to a subset of the observations. Also,
384 each tree in a random forest is fit for a different subset of the predictor variables. For
385 numerical dependent variables, a typical size of the subset of predictors that is considered
386 for each tree is the number of predictors divided by 3 (Hastie et al., 2001; Strobl et al.,
387 2009). The relatively small size of the subset of considered predictors ensures that the
388 trees in a random forest are not too similar.

389 *5.2.2. Prediction and performance*

390 The prediction of a random forest model is defined as the average prediction of the individ-
391 ual trees for the out-of-bag observations (i.e. out-of-bag predictions). The performance
392 of a random forest is evaluated by comparing the average of the out-of-bag predictions
393 with the observed data. The average out-of-bag prediction has less variance and thus
394 suffers less from overfitting when the predictions of individual trees are less correlated.
395 Both bagging and random predictor subset selection ensure that the predictions of the
396 individual trees in a random forest are not too similar. Unlike individual decision trees,
397 random forests therefore tend not to overfit the data and have excellent generalization
398 performance.

399 The interpretation of the results from a random forest are based on a measure of
400 variable importance. Different measures of variable importance exist. The measure we
401 use here is based on permutation tests (Breiman, 2001). To establish the importance of
402 a predictor, the values for that predictor are randomly permuted. The accuracy of the
403 out-of-bag predictions for the permuted predictor is then compared with the accuracy
404 of the out-of-bag predictions for the original predictor. A predictor is regarded to be
405 more important, the greater the difference in prediction accuracy between the original
406 predictor and the permuted predictor.

407 *5.2.3. Predictors and parameters for random forests*

408 The KIEL data set contains a number of categorical predictors. The **ranger** package is
409 able to handle categorical variables, while the **glmnet** package (Friedman et al., 2018)
410 that we will use below to illustrate regularized regression models is not. To be able to
411 compare the variable importances of the random forest with the coefficients in regularized
412 regression on a fair basis, we converted the categorical predictors in the data to numerical
413 variables using *one-hot encoding* that converts the categorical predictors in the KIEL
414 corpus to numerical variables.

415 To understand how one-hot encoding works, consider the categorical predictor **Stress**.
416 **Stress** has three levels: **Primary**, **Secondary** and **None**. We encoded the information
417 in the categorical variable **Stress** in two numerical predictors: **StressPrimary** and
418 **StressSecondary**. **StressPrimary** was set to 1 for words with primary stress and to
419 0 otherwise. Similarly, **StressSecondary** was set to 1 for words with secondary stress
420 and to 0 otherwise. The information for the third level of the categorical variable **Stress**,
421 **None**, is implicitly encoded in **StressPrimary** and **StressSecondary**. Whenever both
422 **StressPrimary** and **StressSecondary** are zero, the word has no stress. We applied
423 one-hot encoding to all categorical predictors in the KIEL data set. (In linear regression
424 modeling, R’s default for categorical predictors, treatment coding, automatically adds
425 such one-hot encoded predictors for factorial predictors to the model matrix.)

426 A crucial parameter in the **ranger()** function is **num.trees**, which determines the
427 number of decision trees that should be fit. The **caret** package (Kuhn, 2018) for R
428 provides grid search functionality for a large number of predictive models, which helps
429 the user tune model parameters. To determine an appropriate value of **num.trees**, we
430 fit a series of random forests with an increasing number of trees to the KIEL corpus data
431 using the **train()** function of the **caret** package. We evaluated the prediction accuracy
432 under (10-fold) cross-validation³

³Cross-validation is a technique to assess the accuracy of a model. The data is partitioned into a training set on the basis of which the model is fit and a test set on the basis of which the accuracy of the model is assessed. In 10-fold cross-validation, the model is trained on 90% of the data and tested on the remaining 10% of the data.

for each model with the root mean squared error (RMSE, which is defined as:

$$RMSE = \sqrt{\frac{\sum_{i=1}^n (\hat{y}_i - y_i)^2}{n}} \quad (5)$$

where \hat{y}_i and y_i are the predicted and observed segment duration for observation i , respectively.

The RMSE for different numbers of trees is presented in Figure 3, bottom. The bottom left panel of Figure 3 shows the results of a coarse grid search, with the number of trees ranging from 1 to 11.5 on a log scale (i.e., from 3 to 98,716 on a non-logged scale). The minimal RMSE in the coarse grid search was observed for a value of 6.5 on the log scale (665 trees, RMSE: 0.5395). We then carried out a second grid search, using numbers of trees near the optimal number of trees in the coarse grid search. The results of this fine grid search are presented in the bottom right panel of Figure 3. The minimal RMSE in the fine grid search was observed for 700 trees (RMSE: 0.5394).

It is worth noting that highly similar RMSEs were observed across a wide range of values of `num.trees`. A post-hoc analysis revealed that the RMSE for models with 23 or more trees were not significantly different from the optimal RMSE. Given the fact that random forests tend to not overfit the data, this is a typical pattern of result in a random forest analysis.

5.2.4. Variable importance

Following the results of the grid searches, we ran the final random forest with the `num.trees` parameter set to 700. The parameter for the number of predictors that are considered in each tree, `mtry`, was set to 10. Unscaled permutation-based variable importances were calculated by setting the value of the parameter `importance` to “permutation” (see Nicodemus et al., 2010, for a discussion of the benefits of unscaled variable importances). Default values were used for all other parameters. The RMSE for the out-of-bag predictions of the final model (0.5394) was nearly identical to the RMSE of the same model under cross-validation.

The variable importances for the random forest are presented in Figure 4. The variable with the highest variable importance is the duration of the word (`wDur`), unsurprisingly. The random forest model furthermore indicates that phonological neighborhood density (NDH) and the number of segments of the word (`nSegperWord`) are highly predictive of

461 segment duration as well. To gain further insight into predictor effects, one can plot
 462 the recursive partitioning tree produced by the `rpart()` function (cf. Figure 3, top, but
 463 allowing for greater tree depth).

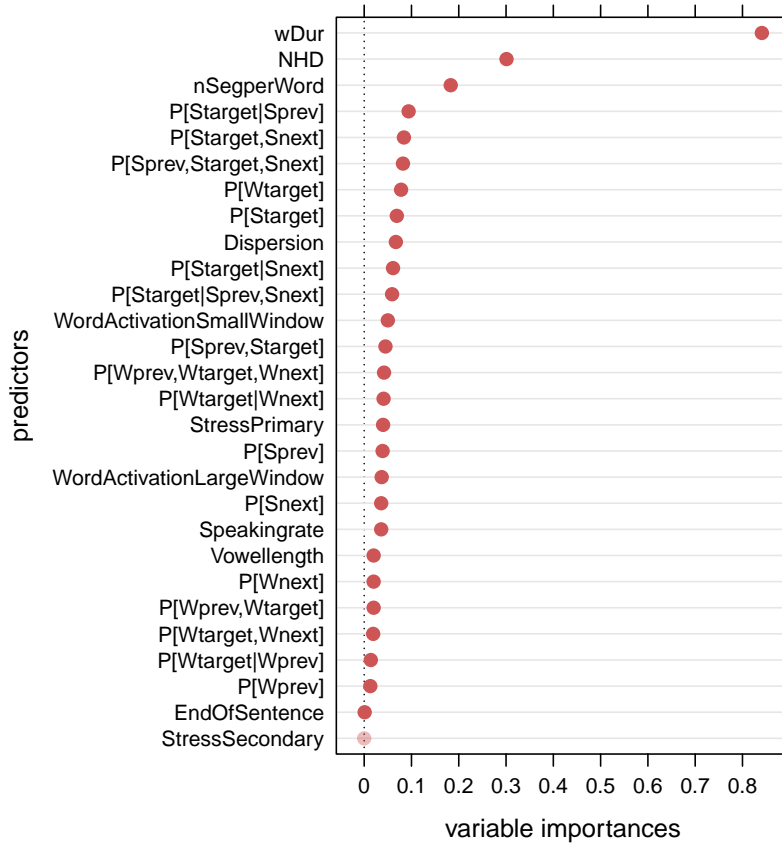


Figure 4: Variable importances for the random forest model. Opaque red dots indicate non-zero variable importances, transparent red dots represent variable importances that are zero.

464 Random forest variable importances provide an excellent assessment of the relative
 465 value of highly correlated predictors. The reason that a predictor gets a chance to show
 466 what it is worth, even though it is highly correlated with an even more powerful predictor,
 467 is that there are trees in the forest in which this more powerful predictor is not included
 468 among the set of predictors for that tree. In a standard recursive partitioning tree that
 469 considers all predictors, each split is based on the most powerful predictor available. In
 470 the forest of trees there are trees where the most powerful predictor is withheld, and
 471 hence the importance of less powerful predictors can be assessed, without the dangers of
 472 suppression or enhancement (see Strobl et al., 2009, for detailed discussion).

473 Random forests provide impressive prediction accuracy. Under cross-validation, a
 474 linear model fit to the segment durations with the 24 numerical covariates explains 50.37%
 475 of the variance in the durations. By contrast, a random forest based on the same set of
 476 predictors explains no less than 70.13% of the variance. As will become apparent below,
 477 none of the other methods for analyzing collinear data comes anywhere close to the
 478 prediction accuracy of the random forest. In the general discussion, we return to this
 479 finding, and discuss its possible theoretical implications.

480 *5.3. Supervised component generalized linear regression*




Figure 5: Simulated data with correlated predictors (left) and the corresponding orthogonalized predictors (right).

481 *5.3.1. Principal components regression and SCGLR*

482 In multivariable regression with k observations and n predictors, an observation i is a
 483 point in a n -dimensional space, whose n axes are set up by the n predictors. When all
 484 predictors are orthogonal, all axes are necessary to define the position of observation i in
 485 this space. When predictors are correlated, there are empty regions in the n -dimensional
 486 space, and a smaller number of axes would suffice to properly locate each datapoint in
 487 a lower-dimensional space. The observations of collinear data sets are points in a space
 488 that, for all practical purposes, has a lower dimensionality than its number of predictors
 489 n .

490 Principal component analysis (Pearson, 1901) is a dimension reduction technique that
491 finds new, orthogonal, axes for the data points, such that the first axis explains the highest
492 proportion of the variance in the space of observations, the second axis explains the next
493 highest proportion of the variance, and so on. For the case of $n = 2$, observations are
494 points on a plane. For the case of $n = 3$, observations are points in a cube. If all
495 points actually lie close to a line in the cube, the first principal component will be a new
496 axis that will be close to all data points. Of the three principal components, the first
497 will explain almost all of the variance. The second and third principal components are
498 superfluous, explaining hardly any variance. Thus, a problem that at first sight appears
499 to be a problem in a three-dimensional space has been reduced to a much simpler problem
500 in a one-dimensional space. This is called dimensionality reduction.

501 Principal components regression is multiple regression that uses principal components
502 derived from the original predictors as regressors. Crucially, not all principal components
503 should be used, otherwise collinearity is back again on the doorstep (Belsley et al., 1980).

504 To make this more concrete, consider Figure 5. The scatter of points in the left
505 panel indicates that predictors A and B are strongly correlated ($r = 0.78$). A principal
506 component analysis rotates the data points anti-clockwise by approximately 130 degrees,
507 resulting in the scatter in the right panel of Figure 5. Most of the variance in the data is
508 now expressed along the horizontal axis, which represents the first principal component
509 (PC1). The remaining variance is found on the vertical axis, which represents the second
510 principal component (PC2). Both principal components are linear combinations of the
511 original A and B axes. The extent to which the old axes are correlated with the new
512 axes is proportional to the so-called loadings of the original variables on the principal
513 components. Principal components are usually entered into a regression analysis simul-
514 taneously. As we have explained above, because they are orthogonal their coefficients
515 will not differ from coefficients obtained in uni-variate models. Principal components
516 regression can be performed using the **ppls** package for R.

517 The goal of a principal components analysis is to reduce the dimensionality of the
518 space in which the observations are points. A commonly used rule of thumb is that the
519 first m components that jointly capture 95% of the variance in the data are retained
520 as new axes (predictors). In a principal components regression, therefore, the c compo-
521 nents that explain very small proportions of the variance are discarded, whereas $k - c$

522 orthogonal predictors are retained as predictors for the response, where k is the original
523 number of predictors. Once a linear model has been estimated for the $k - c$ principal
524 components, the coefficients of the original predictors, given the dimension reduction, can
525 be obtained. The magnitude of these coefficients will be substantially reduced compared
526 to the estimates of a straightforward linear model, whenever the original predictors are
527 substantially collinear.

528 Supervised component generalized linear regression (SCGLR, implemented in the
529 **SCGLR** package, Bry et al. (2013)) builds on the concepts underlying principal compo-
530 nent regression, but the mathematical implementation is substantially different. For the
531 analyst, the important differences are the following.

532 First, SCGLR is designed such that multiple response variables (which can be any of
533 Gaussian, binomial, and Poisson) can be modeled simultaneously. For the KIEL corpus,
534 for instance, we could have included as further predictors the number of segment deletions
535 or syllable durations, the idea being that the predictors for segment duration should also
536 be relevant for understanding segment deletion and syllable duration. In the present
537 survey, space restrictions limit demonstration of this aspect of SCGLR modeling to the
538 supplementary materials.

539 Second, unlike standard principal components regression, SCGLR orthogonalizes not
540 just the predictors, but the predictors and response variables jointly. Whereas principal
541 components regression finds high variance directions in the covariate space, SCGLR sets
542 out to find those directions in the space of the covariates that are optimal for predicting
543 the response variables. Just as in principal components analysis, the components, now
544 called supervised components, are estimated step by step. The first supervised component
545 optimizes a trade-off between the variance it captures in the full variable space (predictors
546 and responses) and the goodness of fit of that component as sole predictor of the response.
547 The second component is selected in the same manner, with the restriction that it has to
548 be orthogonal to the first component. This procedure is repeated until K complementary
549 and mutually independent components are obtained.

550 Third, whereas in principal components regression the number of principal com-
551 ponents to retain is typically based on a rule of thumb, SCGLR implements a cross-
552 validation procedure to determine the optimal number of supervised components.

553 Fourth, SCGLR allows for the possibility that there are predictors that do not need

Table 1: *Coefficients of supervised components and factorial predictors in the SCGLR model.*

predictor	Estimate	Std. Error	t-value	p-value
Intercept	-0.321	0.019	-16.657	0.000
SC1	0.165	0.004	42.804	0.000
SC2	-0.221	0.005	-44.963	0.000
SC3	-0.162	0.005	-30.622	0.000
SC4	-0.096	0.005	-19.055	0.000
SC5	-0.098	0.006	-16.461	0.000
SC6	-0.011	0.006	-1.784	0.074
EndOfSentence	0.122	0.045	2.741	0.006
StressPrimary	0.421	0.022	18.913	0.000
StressSecondary	0.226	0.428	0.528	0.598
Vowellength	0.099	0.021	4.796	0.000

554 to be orthogonalized. For the present data set, such predictors could be the sex and age
 555 of the speaker. Both sex and age are not expected to enter into strong correlations with
 556 the word and segment-bound predictors.

557 5.3.2. Working with SCGLR

558 The steps in an SCGLR analysis are the following. First, the response variables are
 559 selected, and for each response variable, it is determined whether it is Gaussian, binomial,
 560 or Poisson. The single response variable of our working example, `sDur`, is a Gaussian
 561 response.

562 Second, the predictors are grouped into two sets. One set contains the collinear predic-
 563 tors that require orthogonalization, and the other predictors that are not orthogonalized.
 564 For the KIEL data set, the 24 variables laid out in section 3 are assigned to the first set.
 565 The second set comprises the factorial predictors `Stress` (`none`, `primary`, `secondary`),
 566 `EndOfSentence` (`true`, `false`) and `Vowellength` (`long`, `short`).

567 Next, the optimal number K of supervised components needs to be determined. For
 568 this, the **SCGLR** package makes available the function `scglrCrossVal`, which requires
 569 the user to specify the maximum number of components to take into account. We set
 570 this value to 15. As the results of cross-validation may vary from run to run, we carried
 571 out the cross-validation procedure 8 times, and selected the best-supported value, which
 572 turned out to be 6.

573 Finally, the model itself is fit with the `scglr` function, with the parameter K set to
574 6. The model object produced is a list with several components. Of these, the `gamma`
575 component, provides a table of coefficients, together with their standard errors and asso-
576 ciated statistics (see Table 1). The first five supervised components are all well supported
577 as predictors for segment duration, and the same holds for the factorial predictors.

578 The summary of an `scglr` object generates two tables that are essential for the inter-
579 pretation of the supervised components. The `rho` table lists the squared correlations (r^2)
580 of the predictors with the supervised components. The `rho.pred` table provides the same
581 information for the response variables. The information provided by these two tables is
582 merged in Table 2. The first row of this table concerns the response. The greatest r^2 is
583 observed for the second supervised component (SC2). The next largest r^2 is listed for
584 SC1. Thus, in the 6-dimensional space spanned by the 6 SCs, the plane defined by SC2
585 and SC1 is the plane in which the response variable is most strongly represented. This
586 plane is therefore listed in Table 2 as the ‘best plane’. The ‘best value’ is the sum of
587 the r^2 values for the axes of the best plane, and represents the variance in the response
588 captured by the best plane. The remaining rows of Table 2 pertain to the predictors. Like
589 the response, the `Dispersion` measure is most strongly expressed on the plane defined
590 by SC1 and SC2, but for word duration (`wDur`) and speaking rate (`Speakingrate`), the
591 best plane is given by SC3 and SC4.

592 Interpretation of tables such as Table 2 is facilitated by visualization. The `plot` method
593 implemented for `scglr` objects produces correlation plots, examples of which are pre-
594 sented in Figure 6. A correlation plot locates, by means of arrows, variables in the
595 space defined by two (user-selected) supervised components. To avoid visual cluttering,
596 a threshold (represented by a dashed circle) is set such that variables with a best value
597 less than the threshold are not shown. The coordinates of a variable in the plane are
598 the correlations r (the square roots of the values listed in Table 2) of the variable with
599 the pertinent supervised components. The length of a variable’s arrow is, by Pythagoras’
600 theorem, the square root of its best value. Its sign is taken from the correlation between
601 a SC and the original predictor. In Figure 6, the arrows of predictors are presented in
602 black, and that of the response in blue. The threshold was set at 0.5. Measures with best
603 values (arrow lengths) less than 0.5, therefore, are not included in the plots.

604 The left panel of Figure 6 shows that on the SC1 by SC2 plane, neighborhood density

Table 2: *Squared correlations between predictors and supervised components.*

predictor	SC1	SC2	SC3	SC4	SC5	SC6	best plane	best value
sDur	0.330	0.361	0.199	0.059	0.051	0.001	1/2	0.690
Dispersion	0.35	0.40	0.01	0.00	0.06	0.00	1/2	0.751
wDur	0.12	0.01	0.43	0.22	0.02	0.01	3/4	0.656
Speakingrate	0.06	0.02	0.14	0.32	0.10	0.02	3/4	0.454
WordActivation LargeWindow	0.68	0.00	0.05	0.01	0.11	0.00	1/5	0.791
P(Wnext)	0.01	0.04	0.20	0.02	0.01	0.57	3/6	0.777
NHD	0.01	0.69	0.08	0.04	0.00	0.03	2/3	0.764
P(Wtarget)	0.49	0.26	0.01	0.00	0.05	0.00	1/2	0.751
WordActivation SmallWindow	0.66	0.00	0.08	0.04	0.08	0.01	1/5	0.743
nSegperWord	0.01	0.51	0.21	0.07	0.00	0.03	2/3	0.720
P(Sprev, Starget, Snext)	0.56	0.03	0.11	0.09	0.00	0.00	1/3	0.670
P(Wprev)	0.00	0.00	0.13	0.33	0.00	0.28	4/6	0.608
P(Starget Snext)	0.31	0.00	0.01	0.11	0.28	0.00	1/5	0.586
P(Starget Sprev, Snext)	0.36	0.00	0.16	0.00	0.20	0.01	1/5	0.561
P(Starget, Snext)	0.43	0.00	0.01	0.12	0.01	0.00	1/4	0.557
P(Starget)	0.09	0.27	0.06	0.00	0.25	0.00	2/5	0.529
P(Starget Sprev)	0.07	0.18	0.32	0.05	0.06	0.00	2/3	0.499
P(Snext)	0.01	0.03	0.00	0.00	0.47	0.01	2/5	0.497
P(Wprev,Wtarget,Wnext)	0.11	0.06	0.21	0.27	0.01	0.01	3/4	0.483
P(Wtarget Wnext)	0.24	0.01	0.08	0.01	0.04	0.20	1/6	0.440
P(Sprev, Starget)	0.24	0.19	0.09	0.02	0.02	0.00	1/2	0.430
P(Wprev, Wtarget)	0.17	0.01	0.15	0.24	0.02	0.04	1/4	0.406
P(Wtarget Wprev)	0.21	0.00	0.00	0.04	0.01	0.19	1/6	0.399
P(Wtarget, Wnext)	0.18	0.02	0.06	0.10	0.02	0.20	1/6	0.374
P(Sprev)	0.07	0.00	0.11	0.16	0.01	0.00	3/4	0.270

605 (NHD) and word length (nSegperWord) align, with opposite sign, with SC2. Several pre-
606 dictors (P(Sprev, Starget, Snext), WordActivation SmallWindow, WordActivation
607 LargeWindow, P(Starget | Sprev, Snext), P(Starget, Snext)) align with PC1. In
608 the (SC1, SC2) plane, these predictors are orthogonal to word length and neighbor-
609 hood density. A third group of predictors, including P(Wtarget) and P(Starget), are
610 positioned between the axes, with medium correlations on both axes, instead of large
611 correlations with either SC1 or SC2.

612 In the (SC1, SC2) plane, the response, represented by the blue arrow, emerges as

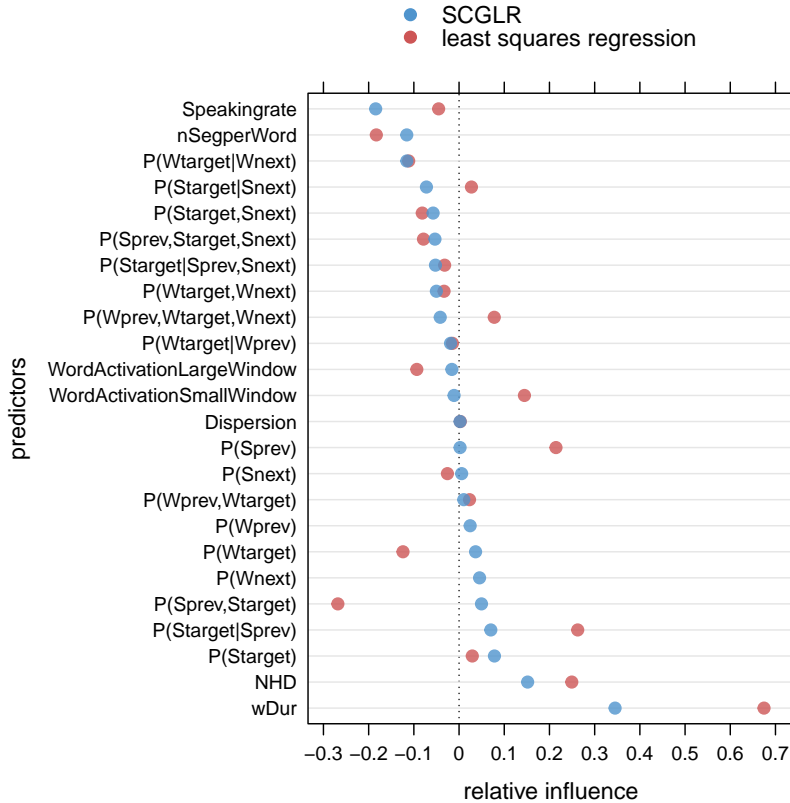


Figure 7: Coefficients for the predictors estimates by the SCGLR (blue) and by a least squares regression model (red). The SCGLR substantially shrinks many of the large coefficients of the regression model towards zero.

630 For assessing the strength of predictors, **SCGLR** makes available a table with the
631 coefficients of the original predictors, which it derives from reduced space of supervised
632 components. These coefficients reflect the cumulative support from all the dimensions
633 of the (reduced) space of orthogonal supervised components. Figure 7 presents these
634 coefficients in blue, together with the corresponding coefficients estimated by a standard
635 linear regression model. Many of the large coefficients of the ordinary least squares re-
636 gression have been shrunk towards zero in the SCGLR. For instance, `WordActivation`
637 `LargeWindow` has a negative coefficient, whereas `WordActivation SmallWindow` has a
638 positive coefficient in ordinary least squares regression. By contrast, the coefficients for
639 both predictors are shrunk towards zero in the SCGLR. Likewise `P(Sprev, Starget)`
640 and `P(Starget | Sprev)` have large coefficients with opposite signs in least squares re-
641 gression, but substantially reduced positive coefficients in the SCGLR model.

642 In other words, the coefficients of the linear model have undergone ‘regularization’: the

643 adverse effects of enhancement have been removed. This will be explained in more detail
644 below in Section 5.4.1. Estimates of uncertainty about the coefficients of the predictors
645 are not available after regularization, however. It is only for the supervised components
646 that standard errors and p-values can be derived.

647 *5.3.3. Advantages and disadvantages of SCGLR*

648 The squared correlation of the model predictions with the response are 0.428 for SCGLR
649 and 0.504 for the standard regression model. When prediction accuracy is of primary
650 importance, SCGLR is therefore a suboptimal choice compared to random forests.

651 What SCGLR does provide is insight into the magnitude and sign of the shrunk
652 predictors. Here, it offers an important advantage over principal component regression.
653 Recall that in contrast to principal components regression, which is designed to find high
654 variance directions in the space of the predictors, SCGLR aims to find dimensions that
655 are optimal for predicting the response. These different design principles enable SCGLR
656 to better distinguish which of a set of correlated predictors are actually predictive for the
657 response. We illustrate this for two highly correlated standard normal predictors, A and
658 B and a dependent variable Y , for which the correlation between A and Y , r_{AY} , is 0.5, and
659 further $r_{BY} = 0$, and $r_{AB} = 0.8$. Analogous to the cases illustrated in Figure 1, a standard
660 regression model will result enhancement, estimating a slope of -1.1 for B even though B
661 is uncorrelated with Y . Orthogonalization with principal components analysis results in
662 one predictor, the first principal component, that has loadings of 0.71 with both A and
663 B . Back-transformed coefficients using the `pcr` function from the `pls` package (Mevik
664 et al., 2018) are 0.14 for both A and B . In other words, the PCA regression does not
665 detect that B is not predictive for Y . However, SCGLR performs much better, with
666 back-transformed coefficients for A and B of 0.42 and 0.05 respectively, a much improved
667 approximation of the actual correlations 0.5 and 0.

668 *5.4. Regression with the elastic net*

669 The elastic net (Zou and Hastie, 2005) is a regression technique that addresses collinear-
670 ity by penalizing overly large β estimates. In this way and unlike in SCGLR, highly
671 collinear predictors may be pruned completely from the data. The elastic net combines
672 the ideas behind two other regularization techniques: the lasso (Tibshirani, 1996) and

673 ridge regression (also known as Tikhonov regularization; Hoerl, 1962; Hoerl and Kennard,
674 1970a,b).

675 Both ridge regression and the lasso penalize non-zero β coefficients in an attempt
676 to improve generalization performance. Ridge regression shrinks non-zero β coefficients
677 towards zero, but never to exactly zero. By contrast, the lasso shrinks β coefficients of
678 variables with limited predictor power to exactly zero. The lasso, therefore, allows for
679 the selection of a set of the most predictive variables. The selection of a set of highly
680 predictive variables is referred to in the machine learning and data mining literature as
681 variable selection, predictor selection, or feature selection.

682 5.4.1. Regularization

For a proper understanding of regularized regression it is important to understand how
 β coefficients are estimated in standard linear regression. Standard linear regression
models are least squares regression models, which minimize the sum of the squares of the
residuals, commonly referred to as the residual sum of squares (henceforth RSS). The RSS
is defined as:

$$\text{RSS} = \sum_{i=1}^n \left(y_i - \left(\beta_0 + \sum_{j=1}^p \beta_j x_{ij} \right) \right)^2, \quad (6)$$

683 where n is the number of observations, p is the number of predictors, y is the response
684 variable, and x_{ij} is the value of predictor j for observation i . The term $y_i - \beta_0 + \sum_{j=1}^p \beta_j x_{ij}$
685 represents the difference between the predicted and the observed values (equivalent to ϵ
686 in Equation 3).

The RSS is small when the squared differences between the observed values (y) and
the predicted values ($\beta_0 + \sum_{j=1}^p \beta_j x_{ij}$) are small. Minimization of the RSS results in a
high-quality fit to the data the model was fit to, but at the cost of suppression and
enhancement for collinear data. Regularized regression, instead of minimizing the RSS,
minimizes the RSS plus a penalty term that makes it costly to have large or many non-zero
 β coefficients. The term that is minimized in the elastic net is:

$$\text{RSS} + \lambda \sum_{j=1}^p \left((1 - \alpha) \beta_j^2 + \alpha |\beta_j| \right). \quad (7)$$

687 The parameter λ determines the strength of the penalty imposed on non-zero β coeffi-
688 cients. As can be seen in Equation 7, both the absolute values of the coefficients ($|\beta|$)
689 and the squared values of the coefficients (β^2) are penalized. The relative weight of the

690 penalties on the absolute values of the coefficients and the squared values of the coeffi-
691 cients is set by the parameter α . The setting of α determines the number of non-zero
692 coefficients in the model, with higher values of α leading to fewer non-zero coefficients.
693 The parameter α thus modulates the extent to which variable selection is performed.
694 When $\alpha = 1$, the model imposes the lasso penalty, and when $\alpha = 0$, the ridge penalty is
695 used.

696 *5.4.2. Data preparation*

697 An implementation of the elastic net for R is available in the **glmnet** package (Friedman
698 et al., 2010, 2018). Before we can run an elastic net model on the KIEL corpus data,
699 we need to prepare the data for analysis with this package. The **glmnet** package does
700 not support categorical predictors. We therefore converted categorical predictors in the
701 KIEL corpus to numerical variables using one-hot encoding, as we did for the random
702 forest analysis in section 5.2.

703 Estimates of the β coefficients are sensitive to the scale of predictors. A change in
704 the scale of a predictor leads to an equivalent change in the scale of the β estimate, but
705 does not influence the RSS of a regression model. By contrast, since the penalty term
706 in regularized regression models takes into account β coefficients, it is sensitive to the
707 scale of predictors. The sensitivity of the penalty term to the scale of predictors has seri-
708 ous consequences for the estimation of the β coefficients in regularized regression models
709 because coefficients for predictors with larger scales are penalized more heavily than co-
710 efficients for predictors with smaller scales. As a result, regularized regression models are
711 biased towards predictors with smaller scales. To prevent regularized regression models
712 from being biased towards predictors with smaller scales, the predictors should be on the
713 same scale. One way to ensure that predictors are on the same scale is standardization,
714 which is enabled by default in the **glmnet** function.

715 *5.4.3. Estimation of parameters*

716 Optimal values of α and λ can be obtained with a grid search. Given a value for α , the
717 **glmnet** function will select an optimal value for λ . By letting α range over a sequence
718 of values between 0 and 1, the optimal values of α and λ can be found. To avoid
719 overfitting, we made use of 10-fold cross-validation, using the **cv.glmnet()** function with
720 the number of folds n set to 10, and using the mean squared error (MSE), i.e. the average

721 of the squared differences between the model predictions and the observed data, as an
722 index of generalization performance. The MSEs reported below are average values of the
723 mean squared error across the 10 folds.

724 Figure 8, top left, shows the cross-validation performance of the elastic net model for
725 different values of α . For each α , the cross-validation score of the best model across 100
726 values of λ is presented. The MSE of the best model for $\alpha = 0$ (which amounts to the
727 ridge penalty), for instance, is 0.484. Error bars represent one standard error confidence
728 intervals. As performance of the elastic net is highly similar for different values of the
729 α , apparently, for the KIEL data set, the influence of the balance between the squared
730 and absolute values of the coefficients on the performance of the model for unseen data
731 is minimal. Nonetheless, since we have to select a value of α , we chose $\alpha = 0.7$, as this
732 value yielded the lowest MSE of 0.480.

733 The center left panel of Figure 8 demonstrates how the MSE for $\alpha = 0.7$ varies with
734 λ under 10-fold cross-validation. To increase readability, λ values are plotted on the
735 log scale, which increases the relative distance between small values of λ selected by
736 the `cv.glmnet()` function. The cross-validation performance of the elastic net model is
737 optimal for the smallest value of λ that we inspected: $\lambda = 0.000472$ (MSE = 0.480). As λ
738 approaches zero, the contribution of the penalty term to the estimate of the coefficients
739 approaches zero as well. As a consequence, the estimated coefficients for small values of
740 λ approach the least squares estimates of the coefficients. The fact that cross-validation
741 performance of the elastic net is optimal for a very small value of λ indicates that a least
742 squares solution may generalize well for the current data.

743 For reasons of interpretability, we increase λ beyond its optimal value to enforce
744 regularization, as larger values of λ result in a smaller number of non-zero coefficients.
745 The key question is how much predictive accuracy we are willing to sacrifice for a more
746 interpretable model. A common strategy is to choose the largest value of λ for which
747 the MSE is within one standard error of the minimum MSE (Breiman et al., 1984; Hastie
748 et al., 2001). For the current model, this approach would lead to fixing λ at 0.0162 ($\log \lambda$
749 = -4.124). For this value of λ , however, no less than 19 predictors still have non-zero
750 coefficients.

Instead of using the one-standard error rule, we therefore placed a threshold on the
percentage by which we allow the MSE of a model to be higher than the minimum MSE.

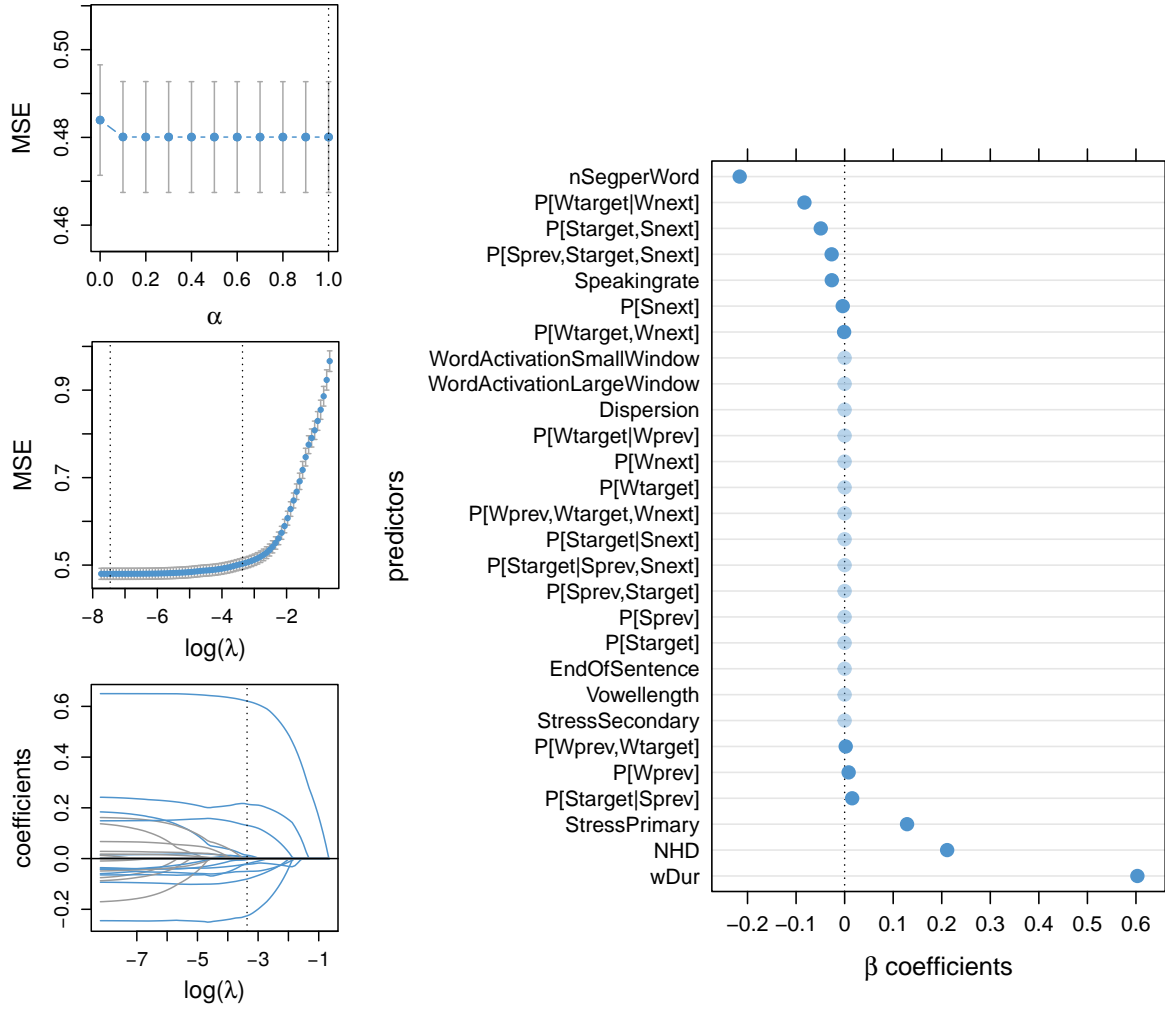


Figure 8: **left:** Results of the elastic net models fit to the segment durations in the KIEL data. Top panel: optimal cross-validation performance for different values of the tuning parameter α . The dashed line indicates the value of α for which the MSE is minimal ($\alpha = 0.7$) across all MSE values. Middle panel: cross-validation performance of the elastic net model with $\alpha = 0.7$ for different (logged) values of the penalty parameter λ . The dashed lines indicate the value of λ for which the MSE is minimal ($\lambda = 0.00047$, $\log \lambda = -7.659$, $\text{MSE} = 0.480$) and the largest value of λ for which the increase in MSE as compared to the MSE for the optimal value of λ is no greater than 5% ($\lambda = 0.0494$, $\log \lambda = -3.007$, $\text{MSE} = 0.504$). Bottom panel: coefficient estimates for the elastic net model with $\alpha = 0.7$ as a function of λ . The dashed line indicates the largest value of λ for which the increase in MSE as compared to the MSE for the optimal value of λ is no greater than 5% ($\lambda = 0.0494$, $\log \lambda = -3.007$). **right:** Coefficient estimates for the elastic net model ($\alpha = 0.7$, $\lambda = 0.0494$). Opaque blue dots indicate non-zero coefficients, transparent blue dots represent coefficients shrunk to zero.

The optimal value for the threshold depends on the relative importance we would like to place on predictive accuracy versus variable selection. Higher threshold values will lead to more variable selection, but less predictive power. We set the threshold to a relatively conservative value of 5%. The greatest value of λ for which the increase in MSE is smaller than or equal to 5% is 0.0494 ($\log \lambda = -3.007$, $\text{MSE} = 0.504$, increase in MSE = 4.94%). This allows us to update the function that regularizes the regression model to:

$$\text{RSS} + 0.0494 \sum_{j=1}^p \left(0.3 * \beta_j^2 + 0.7 * |\beta_j| \right), \quad (8)$$

751 where j is the number of predictors.

752 Figure 8, bottom left, shows how the magnitude of the 24 coefficients is shrunk to-
 753 wards zero as λ is increased. Coefficients that for $\lambda = 0.0494$ ($\log \lambda = -3.007$) and $\alpha = 0.7$
 754 are not completely shrunk to zero are shown in blue, and the coefficients that are pe-
 755 nalized to zero are shown in gray. For extremely small values of λ , the estimates of
 756 the coefficients approximate the least squares estimates of the predictors: no predictor
 757 selection is performed. For very large values of λ , the penalty term is very large and all
 758 coefficients are shrunk to zero. The coefficients for the selected value of λ are located on
 759 the dashed line in Figure 8, bottom left. The right panel of Figure 8 presents the same
 760 shrunk coefficients in a dotplot, non-zero coefficients are represented by opaque blue dots
 761 and coefficients that are zero are represented by transparent blue dots. A total of 15 out
 762 of 28 coefficients were shrunk to zero.

763 Several predictors show the expected pattern of results: Segment durations (`nSegperWord`),
 764 for instance, are substantially shorter for words with more segments (Altmann, 1980) and
 765 greater conditional probability $P(\mathbf{W}_{\text{target}} \mid \mathbf{W}_{\text{next}})$ of the word (Bell et al., 2009). By
 766 contrast, longer word durations (`wDur`), primary word stress (`StressPrimary`, Moon and
 767 Lindblom (1994)), and greater phonological neighborhood density (NHD) lead to longer
 768 segment durations. The direction of NHD is in line with findings by Scarborough (2003)
 769 and Baese-Berk and Goldrick (2009) who report enhancement of a segment’s acoustic
 770 signal in words with greater NHD, but is at odds with recent findings by Gahl and Strand
 771 (2016), who reported shorter word durations for greater NHD.

772 Accurate standard errors for regularized regression models are not available (see Goe-
 773 man, 2010). It is therefore advisable to refrain from reporting p -values for regularized
 774 regression models. Since cross-validated regularized regression models separate the pre-

Table 3: *Estimates of coefficient provided by the elastic net and by a least squares regression model fit to the reduced data set that contains only predictors with non-zero coefficients in the elastic net. Standard errors (S.E.), t-values and p-values are reported for the coefficients estimates of the least squares regression model.*

term	elastic net β	β	S.E.	t-value	p-value
nSegperWord	-0.216	-0.269	0.016	-16.972	< 0.001
P(Wtarget Wnext)	-0.083	-0.108	0.010	-11.324	< 0.001
P(Starget, Snext)	-0.049	-0.057	0.008	-7.292	< 0.001
P(Sprev, Starget, Snext)	-0.027	-0.032	0.010	-3.336	0.001
Speakingrate	-0.026	-0.046	0.007	-6.514	< 0.001
P(Snext)	-0.004	-0.023	0.007	-3.183	0.001
P(Wtarget, Wnext)	-0.001	-0.022	0.007	-3.449	0.001
StressSecondary	0.000	0.276	0.406	0.679	0.497
P(Wprev, Wtarget)	0.002	0.033	0.008	4.176	< 0.001
P(Wprev)	0.008	0.014	0.009	1.467	0.142
P(Starget Sprev)	0.015	0.032	0.007	4.304	< 0.001
StressPrimary	0.129	0.368	0.020	17.980	< 0.001
NHD	0.211	0.237	0.012	19.143	< 0.001
wDur	0.603	0.650	0.011	58.616	< 0.001

775 dictors into effective predictors (with non-zero coefficients) on the one hand, and in-
776 effective predictors (with zero-coefficients) on the other hand, the selection of effective
777 predictors replaces variable selection based on p -values and some (relatively arbitrary)
778 α -level.

779 It is of course possible to fit an unpenalized regression model with only those predictors
780 that have non-zero coefficients in the regularized regression model. The coefficients of
781 such a least squares regression model on the segment durations in the KIEL corpus are
782 presented in Table 3, which also lists the corresponding values given by the elastic net.
783 The two sets of predictors are similar, with the same signs, and a Pearson correlation of
784 $r = 0.913$. There is only one coefficient, that for P(WPrev), that is retained by the elastic
785 net without being significant according to the unpenalized regression. Although for the
786 unpenalized model all variance inflation factors are well below 5, the condition number is
787 still high: 20.22. In this light, it is not surprising that the (absolute) magnitudes of the
788 coefficients of the elastic net are smaller than those of the unpenalized regression, which

Table 4: *Lower triangle of the correlation matrix for the relative influences according to the elastic net, supervised component generalized linear regression (SCGLR), the random forest, and least squares regression.*

	elastic net	SCGLR	least squares
SCGLR	0.565		
least squares	0.934	0.717	
random forest	0.965	0.413	0.861

are, on average, 0.049 and 0.155 respectively. The penalization implemented in the elastic net protects the estimates for the coefficients against collinearity-induced enhancement.

6. Discussion

Random forests, supervised component generalized linear regression, and the elastic net assess collinear data in very different ways. This raises the question of how results obtained with these statistical techniques compare.

To address this question, we need appropriate measures of the relative influence of a predictor. For the random forest analysis, we defined the relative influence of a predictor as its variable importance divided by the sum of the variable importances for all predictors. For the regression models, the relative influence of a predictor was defined as the absolute value of its coefficient divided by the sum of the absolute values of the coefficients of all predictors. For each of the three models, the relative influence of the predictors sums up to 1.

Figure 9 presents the relative influence of the predictors according to the elastic net (blue dots), according to the SCGLR (green dots), according to the random forest (red dots), and according to the least squares regression (yellow dots). The vertical axis shows the predictors in the KIEL data, in descending order of mean relative influence in the four models. As can be seen in Figure 9, the most important predictors have substantial relative influences according to all four modeling techniques. Similarly, the least important predictors have negligible relative influences across models.

Further information about the similarity of the relative influence of the predictors in the different models is presented in Table 4, which lists the correlations between the relative influences of the predictors across the four models. Relative influences of predictors




Figure 9: *Relative influence of the predictors according to the elastic net (blue dots), the SCGLR (green dots), least squares regression (yellow dots), and the random forest (red dots).*

812 are highly similar for the elastic net, least squares regression, and the random forest (all
 813 $r > 0.86$). The results of the random forest thus strongly converge with the results of two
 814 of the three regression techniques. The relative influences in the SCGLR are less similar
 815 to the relative influences in the other models (all $r < 0.72$, minimum $r = 0.413$). In
 816 part, this is due to our decision not to shrink factorial predictors when we fit the SCGLR
 817 model.

818 Prediction accuracy also differs substantially across models: the squared correlation
 819 of predicted and observed segment durations are 0.428 for SCGLR, 0.481 for the elastic
 820 net, 0.504 for the linear model, and 0.701 for the random forest. The low value of R^2 for
 821 SCGLR is unsurprising, as this model works with only 10 parameters, whereas the elastic

822 retains 14 parameters, and the unpenalized linear model has no less than 28 parameters
823 at its disposal (excluding the intercept).

824 The remarkable accuracy of the random forest is due to several factors. First, it is not
825 assumed a-priori that effects of predictors are linear. Second, for SCGLR and the elastic
826 net, we considered main-effect models only, as interactions between numeric predictors
827 are best addressed with the generalized additive model (see Wieling, this issue), and not
828 with the (highly constraining) multiplicative interaction available to the linear model.
829 Conditional inference trees and random forests, however, are able to capture complex
830 interactions involving many predictors. Third, random forests exploit the strengths of all
831 predictors.

832 Thus, the choice of method will depend on the goal of the analysis. If this goal
833 is prediction accuracy, the random forest is by far the best choice. If the goal is to
834 understand the effects of predictors through the sign and magnitude of their slopes in a
835 linear model, the elastic net conveniently weeds out insignificant predictors and provides
836 estimates for the remaining coefficients that are properly shrunk.

837 SCGLR is an informative option when the goal is to better understand the high-
838 dimensional space in which response and predictors are defined, and the joint effect of
839 clusters of predictors on the response is of theoretical interest. Especially for studies
840 in which the predictors are themselves not free of error and are best understood as
841 contributing imperfect probes of the locations of data points in a high-dimensional space,
842 SCGLR comes into its own.

To illustrate this point, consider the relative influence of word duration (`wDur`), num-
ber of segments (`nSegperWord`), and speaking rate (`Speakingrate`) in Figure 9. The
elastic net assigns word duration the greatest relative influence, with number of segments
as runner up. Speaking rate, by contrast, has a small relative influence that is much
reduced compared to that of number of segments. Theoretically, this pattern is puzzling,
as one would expect speaking rate to be the causal factor driving word duration. Fur-
thermore, since the number of segments in a word is a poor man's substitute for word
duration, it is also worrisome that the elastic net values number of segments so much over
speaking rate. The relative influences of these predictors according to SCGLR, by con-
trast, are more intuitive. The relative influence of word duration is muted compared to
unpenalized regression, instead of enhanced, as in the elastic net. Furthermore, speaking

rate is accorded a much higher relative influence that exceeds that of number of segments. Because SCGLR has discovered that these three predictors are strongly represented in the (SC3, SC4) plane (see Figure 6), where they align with the response, it treats them similarly — the coefficients for SC3 and SC4 will give the three predictors the same boost (modulo their individual loadings on the SCs). As a result, their relative influences are more similar to each other. The reason that the elastic net generates very high relative influences for word duration and number of segments is that the penalty

$$\lambda \sum_{j=1}^p \left((1 - \alpha)\beta_j^2 + \alpha|\beta_j| \right).$$

843 in equation (7) can be kept low by substantially penalizing many intermediate coefficients
 844 and only mildly penalizing a few extreme coefficients. Importantly, the way the penalty
 845 is set up has no intrinsic value for linguistic theory, it is just a way to let fewer predictors
 846 do more work in such a way that prediction accuracy is optimized. The result is — indeed
 847 — a model with optimized prediction accuracy, but such a model may not be optimal
 848 from a theoretical perspective.

849 The data set with which we illustrated strategies for the analysis of collinear data
 850 includes information on the speaker, a predictor that within the general framework of
 851 mixed models would be included as a random-effect factor. This raises the question of
 852 how to adapt the three strategies discussed above when random-effect factors need to be
 853 taken into account.

854 Our experience with random forests is that, when participants are included into the
 855 term, partitions are made almost if not totally exclusively on subsets of participants,
 856 typically the largest source of variance. For random-effect factors with many factor levels,
 857 the combinatorics of working through possible partitions typically are too demanding for
 858 conditional inference trees and random forests to be estimable.

859 To our knowledge, there is no version of the elastic net that allows for the inclusion of
 860 random effects as in the linear mixed model (LMM Bates et al., 2014) and the generalized
 861 additive mixed model (GAMM Wood, 2006). It is possible to one-hot encode individ-
 862 ual participants; the mechanism of penalization will ensure that the random effects for
 863 participants will be shrunk.

864 Principal components regression is easy to extend to the LMM and GAMM frame-
 865 works. For instance, a set of collinear predictors bound to items can be orthogonalized

866 using principal components analysis, and pertinent principal components can then be
867 used as predictors for the LMM or GAMM. For fully crossed mixed designs, SCGLR
868 offers the possibility of bringing together subject responses into a multivariate response
869 matrix, to be predicted from the (collinear) set of item-bound predictors. Unlike princi-
870 pal components regression, which orthogonalizes just the space of predictors, SCGLR will
871 search for those directions in the space of the covariates that are optimal for predicting
872 the responses of all of the subjects jointly. The resulting supervised components can, if
873 required, be extracted from the model and used as predictors within a LMM or a GAMM.

874 From the preceding discussion, it will be clear that there are no hard and fast rules
875 for the analysis of multivariate data with substantial collinearity.

876 Each of the statistical methods that we have reviewed has its advantages and dis-
877 advantages, and the choice of a method will depend, to a large extent, on the goals of
878 the analysis. Regression models tend to be well-interpretable, but can be much less ac-
879 curate than random forests. By contrast, random forests tend to provide surprisingly
880 good predictions, but are more like a black box that does not allow inspection of how
881 predictors work together to produce these good predictions. Even when individual trees
882 are inspected, the number of interactions discovered by the tree can be overwhelming.

883 In addition, Important limitations of the regression-based methods is that effects are
884 supposed to be linear, and that interactions of numeric predictors cannot be incorporated
885 in a principled way. The generalized additive model (see Wieling, this volume) does not
886 have these limitations. Unfortunately, the regression methods that we have surveyed are
887 limited to linear (or linearizable) relations between response and predictors.

888 Furthermore, in the nonlinear world, the problem of collinearity resurfaces in the more
889 general form of concurvity. Concurvity can lead to similar problems of interpretation,
890 and it can render model estimates unstable. Concurvity occurs when one smooth term
891 in the model can be approximated by other smooths in the model. This can happen, for
892 instance, if a smooth of time is included together with further smooths for other time-
893 varying covariates. Appendix A provides further information on how concurvity can be
894 assessed, and how one might proceed if substantial concurvity in the model is detected.

895 We conclude with a reflection on the application of statistical analyses. In the context
896 of confirmatory inference for collinear data, with as goal establishing whether a particular
897 covariate is significant, the elastic net seems a good choice. If the covariate is not shrunk

898 to zero, it can be accepted as supported, possibly in combination with further support
899 from a least squares regression that discards all predictors that have been shrunk to zero
900 by the net. For exploratory data analysis, all methods surveyed above are useful. The
901 multiple testing method of Goeman and Solari (2011); Meijer and Goeman (2015), which
902 is designed specifically for exploratory data analysis of collinear data, is an excellent
903 companion to SCGLR.

904 **7. Acknowledgements**

905 This research was partially funded by two grants from the Deutsche Forschungsgemein-
906 schaft (German Research Foundation) (BA 3080/3-1 and 387774888). Our thanks go to
907 Theresa Schmitt for her help with preprocessing the KIEL corpus.

908 **References**

- 909 Adelman, J., Brown, G., and Quesada, J. (2006). Contextual diversity, not word fre-
910 quency, determines word-naming and lexical decision times. Psychological Science,
911 17(9):814.
- 912 Altmann, G. (1980). Prolegomena to menzerath's law. Glottometrika, 2:1–10.
- 913 Aylett, M. and Turk, A. (2004). The smooth signal redundancy hypothesis: A functional
914 explanation for relationships between redundancy, prosodic prominence, and duration
915 in spontaneous speech. Language and Speech, 47(1):31–56.
- 916 Baayen, R. H. (2008). Analysing linguistic data: A practical introduction to statistics.
917 languageR package Version 1.4.1. Cambridge University Press, Cambridge, MA.
- 918 Baayen, R. H., Feldman, L., and Schreuder, R. (2006). Morphological influences on the
919 recognition of monosyllabic monomorphemic words. Journal of Memory and Language,
920 53:496–512.
- 921 Baayen, R. H., Milin, P., and Ramscar, M. (2016). Frequency in lexical processing.
922 Aphasiology, 30(11):1174–1220.
- 923 Baayen, R. H., Milin, P., Đurđević, D. F., Hendrix, P., and Marelli, M. (2011). An
924 amorphous model for morphological processing in visual comprehension based on naive
925 discriminative learning. Psychological review, 118(3):438–481.
- 926 Baayen, R. H., Vasishth, S., Kliegl, R., and Bates, D. (2017). The cave of shadows. ad-
927 dressing the human factor with generalized additive mixed models. Journal of Memory
928 and Language, pages 206–234.
- 929 Baese-Berk, M. and Goldrick, M. (2009). Mechanisms of interaction in speech production.
930 Language and Cognitive Processes, 24:527–554.
- 931 Bates, D., Maechler, M., Bolker, B., and Walker, S. (2014). lme4: Linear mixed-effects
932 models using Eigen and S4.
- 933 Bell, A., Brenier, J. M., Gregory, M., Girand, C., and Jurafsky, D. (2009). Predictability
934 effects on durations of content and function words in conversational English. Journal
935 of Memory and Language, 60(1):92 – 111.
- 936 Belsley, D. (1984). Demeaning conditioning diagnostics through centering.
937 The American Statistician, 38:73–77.
- 938 Belsley, D. A., Kuh, E., and Welsch, R. E. (1980). Regression Diagnostics. Identifying

939 Influential Data and sources of Collinearity. Wiley Series in Probability and Mathe-
940 matical Statistics. Wiley, New York.

941 Box, G. E. P. (1976). Science and statistics. Journal of the American Statistical
942 Association, 71:791–799.

943 Breiman, L. (2001). Random forests. Machine Learning, 45(1):5–32.

944 Breiman, L., Cutler, A., Liaw, A., and Wiener, A. (2018). Package ‘randomForest’.

945 Breiman, L., Friedman, J., Olshen, R. A., and Stone, C. J. (1984). Classification and
946 decision trees. Wadsworth and Brooks, Monterey, CA.

947 Bry, X., Trottier, C., Verron, T., and Mortier, F. (2013). Supervised component gener-
948 alized linear regression using a pls-extension of the fisher scoring algorithm. package
949 version 2.0.3. Journal of Multivariate Analysis, 119:47 – 60.

950 Chatterjee, S. and Hadi, A. (2012a). Regression analysis by example. John Wiley &
951 Sons, New York.

952 Chatterjee, S., Hadi, A., and Price, B. (2000). Regression analysis by example. John
953 Wiley & Sons, New York.

954 Chatterjee, S. and Hadi, A. S. (2012b). Regression Analysis by Example. Fifth Edition.
955 Wiley.

956 Farrar, D. E. and Glauber, R. R. (1967). Multicollinearity in regression analysis: The
957 problem revisited. The Review of Economics and Statistics, 49(1):92–107.

958 Fox, J. and Weisberg, S. (2011). An R Companion to Applied Regression. Car Package
959 Version 2.1-6. Sage, Thousand Oaks CA, second edition.

960 Friedman, J., Hastie, T., and Tibshirani, R. (2010). Regularization paths for generalized
961 linear models via coordinate descent. Journal of Statistical Software, 33(1):1–22.

962 Friedman, J., Hastie, T., Tibshirani, R., Simon, N., Narasimhan, B., and Qian, J. (2018).
963 Package ‘glmnet’, Version 2.0-13.

964 Friedman, L. and Wall, M. (2005). Graphical views of suppression and multicollinearity
965 in multiple linear regression. The American Statistician, 59(2):127–136.

966 Gahl, S. (2008). ”thyme” and ”time” are not homophones. word durations in spontaneous
967 speech. Language, 84(3):474–496.

968 Gahl, S. and Strand, J. (2016). Many neighborhoods: Phonological and perceptual
969 neighborhood density in lexical production and perception. Journal of Memory and
970 Language, 89:162 – 178.

- 971 Goeman, J. J. (2010). L1 penalized estimation in the Cox proportional hazards model.
972 Biometrical Journal, 52(1):70–84.
- 973 Goeman, J. J. and Solari, A. (2011). Multiple testing for exploratory research. Statistical
974 Science, 26(4):584–597.
- 975 Hadi, A. S. (1988). Diagnosing collinearity-influential observations. Computational
976 Statistics and Data Analysis, 7(2):143 – 159.
- 977 Hastie, T., Tibshirani, R., and Friedman, J. (2001). The Elements of Statistical Learning.
978 Springer, New York.
- 979 Hoerl, A. E. (1962). Application of ridge analysis to regression problems. Chemical
980 Engineering Progress, 58:54—59.
- 981 Hoerl, A. E. and Kennard, R. W. (1970a). Ridge regression: Applications to nonorthog-
982 onal problems. Technometrics, 12(1):69–82.
- 983 Hoerl, A. E. and Kennard, R. W. (1970b). Ridge regression: Biased estimation for
984 nonorthogonal problems. Technometrics, 12(1):55–67.
- 985 Hothorn, T., Hornik, K., Strobl, C., and Zeileis, A. (2018a). Package ‘‘party’’.
- 986 Hothorn, T., Seibold, H., and Zeileis, A. (2018b). Package ‘‘partykit’’.
- 987 Jurafsky, D., Bell, A., Gregory, M., and Raymond, W. D. (2000). Probabilistic relations
988 between words: Evidence from reduction in lexical production. In Bybee, J. and Hop-
989 per, P., editors, Frequency and the emergence of linguistic structure. John Benjamins,
990 Amsterdam.
- 991 Keuleers, E., Stevens, M., Mandera, P., and Brysbaert, M. (2015). Word knowledge in the
992 crowd: Measuring vocabulary size and word prevalence in a massive online experiment.
993 The Quarterly Journal of Experimental Psychology, (8):1665–1692.
- 994 Kohler, K. J. (1996). Labelled data bank of spoken standard German – The Kiel Corpus
995 of read/spontaneous speech.
- 996 Kuhn, M. (2018). Package ‘‘caret’’, Version 3.3.
- 997 Meijer, R. J. and Goeman, J. J. (2015). A multiple testing method for hypotheses
998 structured in a directed acyclic graph. Biometrical Journal, 57(1):123–143.
- 999 Mevik, B.-H., Wehrens, R., Liland, K. H., and Hiemstra, P. (2018). Package ‘‘pls’’,
1000 Version 2.6-0.
- 1001 Milin, P., Feldman, L. B., Ramscar, M., Hendrix, P., and Baayen, R. H. (2017). Discrim-
1002 ination in lexical decision. PLOS-one, 12(2):e0171935.

- 1003 Moon, S.-J. and Lindblom, B. (1994). Interaction between duration, context, and speak-
1004 ing style in english stressed vowels. he Journal of the Acoustical Society of America,
1005 96:40–55.
- 1006 Nicodemus, K. K., Malley, J. D., Strobl, C., and Ziegler, A. (2010). The behaviour of
1007 random forest permutation-based variable importance measures under predictor corre-
1008 lation. BMC Bioinformatics, 11(1):110.
- 1009 Pearson, K. (1901). On lines and planes of closest fit to systems of points in space.
1010 Philosophical Magazine, 2(6):559–572.
- 1011 Peters, B. (2003). Die Datenbasis The Kiel Corpus.
- 1012 Priva, U. C. (2015). Informativity affects consonant duration and deletion rates.
1013 Laboratory Phonology, 6(2):243–278.
- 1014 R Core Team (2018). R: A Language and Environment for Statistical Computing, Version
1015 3.3.3. R Foundation for Statistical Computing, Vienna, Austria.
- 1016 Scarborough, R. (2003). Lexical confusability and degree of coarticulation. Annual
1017 Meeting of the Berkeley Linguistics Society, 29(1):367–378.
- 1018 Sheather, S. (2009). A modern approach to regression with R. Springer Science &
1019 Business Media.
- 1020 Strobl, C., Malley, J., and Tutz, G. (2009). An introduction to recursive partitioning:
1021 Rationale, application and characteristics of classification and regression trees, bagging
1022 and random forests. Psychological Methods, 14(4):323–348.
- 1023 Therneau, T., Atkinson, B., and Ripley, B. (2017). rpart: Recursive Partitioning and
1024 Regression Trees. R package version 4.1-11.
- 1025 Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. Journal of the
1026 Royal Statistical Society, 58(1):267–288.
- 1027 Tomaschek, F., Plag, I., Ernestus, M., and Baayen, R. H. (2018). Modeling the duration
1028 of word-final s in English with naive discriminative learning. Manuscript, University
1029 of Siegen/Tübingen/Nijmegen.
- 1030 Tremblay, A. and Tucker, B. V. (2011). The effects of n-gram probabilistic measures on
1031 the recognition and production of four-word sequences. The Mental Lexicon, 6(2):302–
1032 324.
- 1033 Venables, W. N. and Ripley, B. D. (2002). Modern Applied Statistics with S, Version
1034 7.3-45. Springer, New York, fourth edition. ISBN 0-387-95457-0.

- 1035 Wei, T., Simo, V., Levy, M., Yihui, X., Jin, Y., and Zemla, J. (2017). Package ‘‘corrplot’’,
1036 Version 0.84.
- 1037 Wood, S. N. (2006). Generalized additive models: an introduction with R. Chapman
1038 and Hall/CRC, Boca Raton, Florida, U. S. A.
- 1039 Wright, M. N. and Ziegler, A. (2017). ranger: A fast implementation of random forests for
1040 high dimensional data in C++ and R. version 0.10.1. Journal of Statistical Software,
1041 77(1):1–17.
- 1042 Wurm, L. H. and Fisicaro, S. A. (2014). What residualizing predictors in regression
1043 analyses does (and what it does not do). Journal of Memory and Language, 72:37–48.
- 1044 York, R. (2012). Residualization is not the answer: Rethinking how to address multi-
1045 collinearity. Social Science Research, 41(6):1379 – 1386.
- 1046 Zou, H. and Hastie, T. (2005). Regularization and variable selection via the elastic net.
1047 Journal of the Royal Statistical Society, 67(2):301–320.
- 1048 Zuur, A. F., Ieno, E. N., and Elphick, C. S. (2010). A protocol for data exploration to
1049 avoid common statistical problems. Methods in Ecology and Evolution, 1(1):3–14.

1050 **Appendix A: Concurvity**

1051 For the analysis techniques in the main text we assumed that the effects of covariates
1052 are linear. To relax the linearity assumption, we exchange (regularized) regression mod-
1053 eling by regression with the generalized additive model (see Wieling, this volume, for an
1054 introduction). In the nonlinear world, the problem of collinearity resurfaces in the more
1055 general form of concurvity. Concurvity can lead to similar problems of interpretation,
1056 and can make model estimates to some extent unstable. Concurvity occurs when one
1057 smooth term in the model can be approximated closely by other smooths in the model.

1058 The **mgcv** package provides a function `concurvity`, that calculates several related
1059 indices that all range between 0 and 1. The closer the concurvity index for a smooth
1060 is to 1, the greater the risk of a lack of identifiability of a clear estimate. The indices
1061 are all based on a decomposition of a given smooth f into two parts, a part u that is
1062 unique to that smooth’s space, and a part g that lies completely in the space of one or
1063 more other smooths. The indices evaluate how g compares to f . In what follows, we
1064 consider the index that is the ratio of the squared Euclidean lengths of the vectors of f

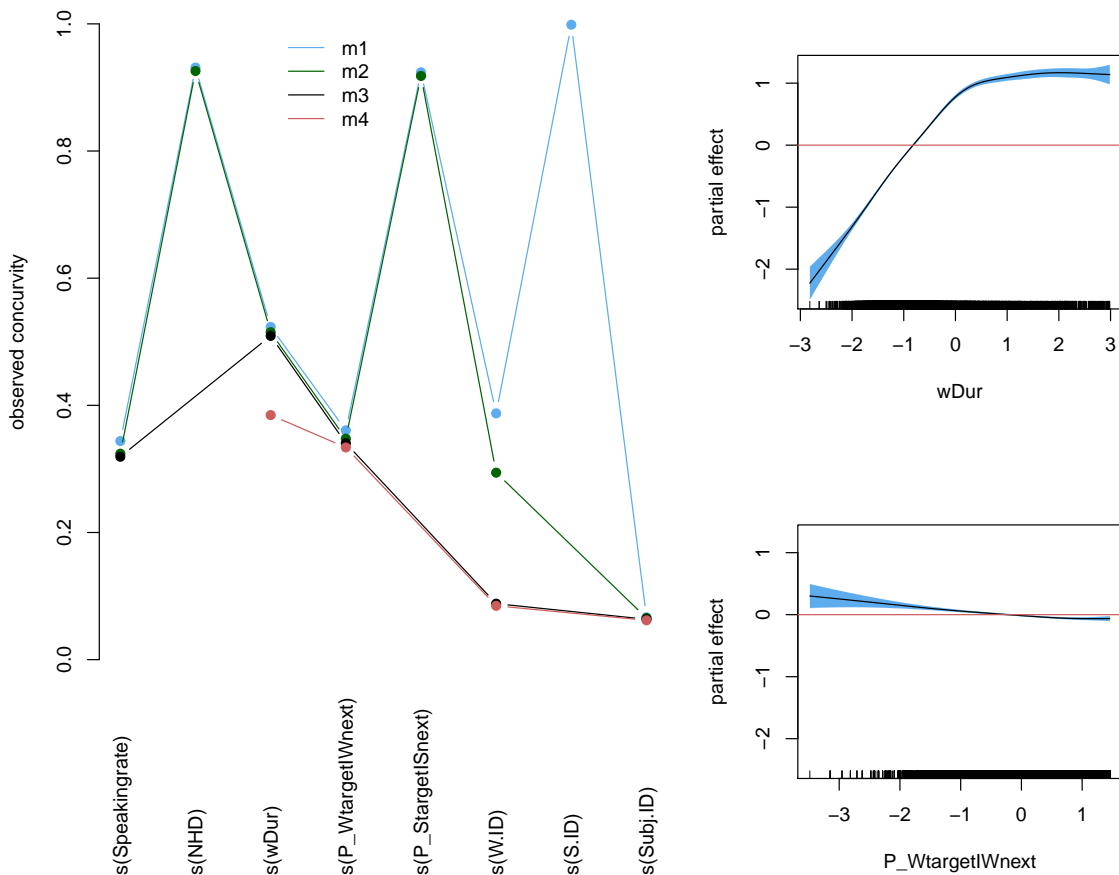


Figure 10: Observed concavity for models m1 (blue), m2 (dark green), m3 (black) and m4 (red), left panel, and the nonlinear effects of word duration and $P(W_{target} | W_{next})$ (right panel).

1065 and g when evaluated at the observed values of the covariates. This measure is possibly
 1066 somewhat over-optimistic, for more pessimistic measures, the reader is referred to the
 1067 documentation of the `concurvity` function.

1068 We illustrate how concavity can be diagnosed and addressed by fitting a generalized
 1069 additive mixed model to the segment durations in the KIEL corpus. We include ran-
 1070 dom intercepts for `speaker`, `word`, and `segment`, and the top seven best predictors that
 1071 emerged from the analyses in the main text: `Speakingrate`, `nSegperWord`, `NHD`, `wDur`,
 1072 `Stress`, $P(W_{target} | W_{next})$, and $P(S_{target} | S_{next})$ (see Figure 9). With the ex-
 1073 ception of `nSegperWord`, all numerical variables were modeled with thin plate regression
 1074 spline smooths. The left panel of Figure 10 presents four GAM models with different sets
 1075 of predictors. Model m1 (blue) includes all predictors, whereas model m4 (red) includes
 1076 only two random effect factors, speaker and word, and only two smooths terms (`wDur`

1077 and $P(\text{Wtarget} \mid \text{Wnext})$). Models **m2** and **m3** are intermediate between **m1** and **m4** with
1078 respect to the predictors included. The left panel of Figure 10 presents the observed
1079 concurvity for each model. For the full model (**m1**), the random intercepts for segment
1080 emerge as completely unidentifiable. This model is clearly overspecified. But the neigh-
1081 borhood density measure (**NHD**) and the probability $P(\text{Starget} \mid \text{Snext})$ also are not well
1082 identifiable — they contribute little that is not already contributed by other predictors.
1083 Model **m2** removes the by-segment random intercepts, but this does little to alleviate the
1084 problems with **NHD** and $P(\text{Starget} \mid \text{Snext})$. Model **m3** removes these two predictors
1085 from the model specification, and model **m4** removes **Speakingrate**, which was not well
1086 supported, thereby reducing the concurvity for **wDur** (which is strongly correlated with
1087 speaking rate). The right panel presents the nonlinear effects of **wDur** and $P(\text{Wtarget} \mid$
1088 $\text{Wnext})$ in model **m4**; both predictors show muted effects for higher values, especially so
1089 for word duration.

1090 In summary, when effects are nonlinear, concurvity may make it impossible to identify
1091 the unique contributions of predictors, even when model summaries suggest predictors
1092 are well supported. The problem is not that the requested model cannot be fit, or that
1093 the requested model does not improve on simpler models. Rather, the problem is that
1094 especially in the nonlinear world, the unique contribution of strongly correlated predictors
1095 will often not be separable. In this case, to further understanding without overfitting the
1096 data, while at the same time complying with Occam’s razor, it is best to keep the model
1097 simple by removing predictors with high concurvity indices.

1098 In the context of confirmatory data analysis where model **m1** was the planned model,
1099 the removal of unidentifiable predictors would be part of model criticism, with as aim to
1100 obtain more reliable estimates of the effects (see Baayen et al. (2017), for discussion of
1101 the importance of model criticism in the context of confirmatory data analysis).