Linguistic Correlates of Acoustic Duration in English Triconstituent Compounds

Annika Schebesta^{*1} and Gero Kunter¹

¹*Universität Siegen*

Abstract

This paper investigates the effect of a range of variables on the acoustic duration of constituents in left- and right-branching NNN compounds derived from a corpus of spoken English (Boston University Radio Speech Corpus, Ostendorf et al. 1997). Variables of interest include speaker-dependent as well as phonological, morphological and lexical factors. The analysis reveals that most variables affect constituent durations as expected, and only few predictors do not yield any effect on the acoustic signal. Furthermore, we detected a complex interplay of the morphological structure of NNN compounds and the two involved bigram frequencies. For instance, the duration of N2 in left-branching compounds is affected by the frequency of N2N3 even though these two constituents do not form a morphological unit in this type of NNN compound. This interplay may be interpreted as a way of resolving potential conflicts between the frequency of adjacent constituents and the morphological structure: In such an instance, speakers appear to use acoustic duration to signal the branching direction of the triconstituent compound.

Keywords: English compounds, acoustic duration, morphological structure

1 Introduction

The acoustic duration of words in spoken utterances is known to be affected by a wide range of variables. Some of these variables may be considered extra-linguistic in the sense that they do not relate to the language system itself but rather to the speakers that use that language, or to the context in which the language is used. An obvious example for such an extra-linguistic variable is the speech rate of a speaker, typically expressed as the number of linguistic segments (words, syllables, phonemes) per time unit. While the speech rate may be used to convey pragmatic content (cf. Uhmann 1992), it is usually not used to express linguistic contrasts.

Other variables affecting acoustic duration are clearly more linguistic, such as the number of segments in the phonological representation of a word that will co-determine how long it

^{*}corresponding author: Annika.Schebesta@uni-siegen.de, Universität Siegen, Seminar für Anglistik, Adolf-Reichwein-Straße 2, 57068 Siegen, Germany

takes to pronounce this word. Somewhat less obvious, the presence of boundaries has been found to affect the duration of linguistic units in their vicinity. For instance, words produced before a prosodic boundary tend to be longer than elsewhere (e.g. Turk and Shattuck-Hufnagel 2007), and the presence of morphological boundaries affects the duration of the base in morphologically complex words (e.g. Lehiste 1972).

Previous research has investigated lexical factors that reflect how speakers and listeners process linguistic material. By now it is well established that the frequency of occurrence of a word not only facilitates lexical processing, but also has a reductive effect in production: highly frequent words are generally shorter than infrequent words (e.g. Wright 1979; Pluymaekers et al. 2005; Bell et al. 2009). Also, it has been shown that predictability affects acoustic duration: A word that occurs in a linguistic context where its occurrence is highly predictable is generally produced shorter than in a less predictable linguistic context (e.g. Bell et al. 2003; Aylett and Turk 2006; Bell et al. 2009).

This incomplete overview illustrates that the acoustic duration of words may be influenced by variables coming from very different domains. Crucially, while the effect of most of these variables has been empirically investigated in previous research, these investigations usually focus on only one single variable or a small set of variables, and it is not always clear which variable contributes to acoustic duration to what extent.

One of the purposes of this paper is to disentangle the interplay between several of the above mentioned variables so that the individual contribution of each becomes clearer. We will focus on variables that pertain to the phonological and morphological features of the items under investigation, as well as on variables pertaining to their representation and interconnections in the mental lexicon. The testbed that we will use is a particular type of morphologically complex words: English NNN compounds, i.e. English compounds that combine three nominal lexemes into a single lexical unit. This paper investigates the acoustic duration of the three individual constituents of such complex words and the linguistic and extra-linguistic variables that affect it. English compounds are well-suited for this type of investigation due to their flexible nature. In contrast to, for instance, derivation as another type of word-formation, compounding has the advantage that the involved linguistic elements are lexemes in their own rights. Thus, this word-formation process can be used to combine lexical words with very diverse phonological make-ups and different lexical features into morphologically complex words, and these words can be found in different prosodic contexts.

Due to the binary structure of compounds (see, e.g. Bauer 2017, ch. 3), triconstituent NNN compounds may be analysed as the result of a two-level application of composition. On the lower level, two nominal lexemes are combined to form an NN compound, and on the higher level, this compound (the ‘embedded compound’) is combined with a third nominal lexeme (the ‘free constituent’) so that the NNN compound is formed. The position of the embedded compound within the NNN compound determines the branching direction, as indicated in (1). Triconstituent compounds such as [*health care*] *law* are left-branching because the embedded compound is in the left position, whereas a compound such as *corner* [*drug store*], with the embedded compound in the right position, is right-branching. Our analysis

includes compounds of both branching directions, which allows us to systematically investigate whether the internal structure of complex words is a morphological factor influencing the acoustic duration of compound constituents.

(1) Branching direction in NNN compounds

The remainder of this paper is structured as follows. Section 2 presents a detailed survey of previous research on the linguistic correlates of acoustic duration. In section 3, we describe the methodology that we used to investigate whether these linguistic correlates allow us to determine the acoustic duration of the constituents in a set of NNN compounds gathered from a corpus of spoken American English. The findings of this analysis are presented in section 4 followed by a discussion and concluding comments in section 5, which put the results into a wider perspective.

2 Factors affecting the acoustic duration of compound constituents

This section provides an overview of factors that have been shown in the literature to have an impact on the acoustic duration of words in general, and of compound constituents in particular. It comprises one extra-linguistic factor as well as several linguistic variables that either pertain to the structure of the mental lexicon or describe the phonological and morphological structure of the involved lexemes.

Extra-linguistic factors Perhaps the most obvious variable that affects the acoustic duration of a spoken word is an extra-linguistic variable: the overall speech rate of the speaker. Speech rate has been found to affect vowel duration (Johnson et al. 1993, Smith 2002) and segment reduction (Raymond et al. 2006) in such a way that a higher speech rate leads to an overall shorter duration of utterances, while a lower speech rate results in longer durations. Speech rate is usually defined as the number of linguistic elements per time interval, for example words, syllables, or segments per second. Different speakers will have different speech rates, and the speech rate of a given speaker is variable as well (Uhmann 1992). Thus, the same word will have different durations for different speakers, and durational variation of the same word is also possible if speakers vary their speech rate.

Lexical factors For the purpose of this paper, we consider all variables that pertain to the storage of lexical items and previous experience with those items as ‘lexical variables’. Lexical frequency is such a correlate of acoustic duration which plays a role in this paper. For example, Wright (1979), Pluymaekers et al. (2005), Aylett and Turk (2006), and Bell et al. (2009) report that words with higher lexical frequencies are typically shorter than corresponding low-frequency words. In Gahl (2008)’s analysis of lexical frequency effects on homophonous words, she reveals that frequent homophones are pronounced shorter than their infrequent counterparts. Beyond single words, Arnon and Cohen Priva (2013) find that the frequency of strings of words occurring together can affect the acoustic duration in speech production, as well. Their results from corpus data and experimental data reveal that higher n-gram frequencies lead to shorter word durations. Summing up, there is strong empirical evidence that high-frequency words – or strings of words – have comparatively shorter durations than low-frequency words or strings of words. Hence, three different types of lexical frequency are of interest in NNN compounds: (a) the constituent frequency, i.e. the frequency of each constituent alone, (b) the bigram frequency of two neighboring compound constituents, and (c) the trigram frequency, i.e. the frequency of the whole compound. In this study, lexical frequency refers to the token frequency of the involved constituents, bigrams and trigrams in the DVD version of the *Corpus of Contemporary American English* (henceforth COCA, Davies 08), covering the years 1990–2012. High token frequencies of a constituent, of a bigram and of a trigram in the reference corpus are expected to have a shortening effect on the acoustic duration of the involved constituents.

The number of phonologically similar words is another measure that is known to affect word durations. In Gahl et al. (2012) it has been shown that the phonological neighborhood density, i.e. the number of words in the lexicon of a speaker that differ from the target word in only one phoneme, affects both word recognition and speech production. They find that target words in a dense neighborhood with a high number of phonological neighbors have shorter durations than target words in a sparse neighborhood.

Words without phonological neighbors, so-called lexical hermits, cannot be accounted for with phonological neighborhood density. Instead, Suárez et al. (2011) suggest that the phonological Levenshtein distance (PLD20) is more appropriate in such cases. This metric computes the mean number of segments that need to be changed in the target word (i.e. deleted, substituted or added) to yield the 20 nearest neighbors. Suárez et al. (2011) show that a low PLD value indicates a high number of phonologically similar lexemes surrounding the target word, and that this target word is more predictable than words with higher PLD values. The effect of PLD20 should be similar to that of phonological neighborhood density: a low PLD value should lead to shorter acoustic duration than a high PLD value. However, this prediction has not been empirically tested yet.

Previous research has shown that the predictability of linguistic material affects acoustic duration. For instance, Gregory et al. (1999) and Jurafsky et al. (2001) show that higher lexical probability leads to reduced word forms, as well as Bell et al. (2003), who find vowel reduction and overall shorter durations in more predictable words (see also Seyfarth 2014 for

an overview of probabilistic reduction).

Phonological factors The number of phonemes is a fairly obvious phonological factor determining word duration: other things being equal, a word that contains more phonemes will be longer than a word with fewer phonemes. In addition, the average duration of the phonemes interacts with the number of syllables in a word. In their review of the effect of word boundaries on acoustic duration, Turk and Shattuck-Hufnagel (2000) illustrate that polysyllabic shortening, a term coined by Lehiste (1972), is a frequent phenomenon. In polysyllabic shortening, the duration of a base word (e.g. *shake*) decreases when additional syllables are added to it, either by suffixation (*shake+ing*) or in a compound (*shakedown*).

The correlation of duration and accentuation has been examined extensively in previous research. The presence of a pitch accent on a syllable has a lengthening effect on the syllable and, by extension, on the word containing it (Turk and Sawusch 1996, Turk and White 1999, de Jong 2004, Kunter 2011, Morrill 2012). The findings reported in Turk and White (1999) suggest that this effect is not constrained to the accented syllable itself, but that there is rightward spreading of accentuation into the material following an accented syllable.

With regard to English compounds in particular, it has been argued in Kunter (2011) that the two possible prominence patterns in English NN compounds ('left-prominent' as in *watch maker* and 'right-prominent' as in *apple pie*, which corresponds to what has traditionally been called 'compound stress' and 'phrasal stress', respectively, cf. Liberman and Prince 1977, Giegerich 2004) are phonologically realized by different accentuation patterns. Accordingly, the first constituent in left-prominent compounds was found in Kunter (2011) to be longer than the first constituent in right-prominent compounds, a finding similarly reported in Farnetani et al. (1988).

Building on these findings, it follows that the duration of the constituents in triconstituent compounds will also be affected by their prominence pattern. The branching direction of the triconstituent compound has been proposed as one of the most important determinants of the compound's prominence pattern. For instance, the Lexical Category Prominence Rule introduced by Liberman and Prince (1977) states that the constituent of a compound can only be prominent if it branches. Accordingly, this rule predicts that in left-branching triconstituent compounds [N1 N2] N3 (e.g. [*law degree*] *requirement*), the most prominent constituent will be the first one (N1, *law*), while in right-branching compounds N1 [N2 N3] (e.g. *degree* [*language requirement*]), the second constituent (N2, *language*) will be most prominent. These predictions were partially confirmed by Kösling and Plag (2009)'s acoustic analysis of tri-constituent compounds sampled from a spoken corpus. However, the study revealed a large number of compounds with aberrant prominence patterns (e.g. left-branching compounds in which the second constituent was most prominent or right-branching compounds in which the third constituent was most prominent). These findings gave rise to an alternative to the Lexical Category Prominence Rule, namely the Embedded Prominence Hypothesis (Kösling et al., 2013). According to this hypothesis, the most prominent constituent in an NNN compound is the one constituent that is most prominent in the embedded compound. Thus, in

left-branching [N1 N2] N3, the most prominent constituent will be either N1 or N2, depending on whether the embedded compound [N1 N2] is left- or right-prominent. In right-branching compounds N1 [N2 N3], the hypothesis predicts that if the embedded compound [N2 N3] is left-prominent, N2 will be most prominent, otherwise most prominence will fall on N3. These predictions find support in the acoustic analysis of Kösling et al. (2013)’s experimental data: their modeling of pitch contours across different speakers and compounds reveals an accentuation pattern that is mostly compatible with the Embedded Prominence Hypothesis. What all this illustrates is that the accentuation pattern in NNN compounds, affected both by branching direction and the prominence pattern of the embedded compound, will affect the acoustic duration of the compound constituents.

Prosodic boundaries have been shown to have an impact on the phonological structure and the phonetic signal of words (e.g. Selkirk 1980, Hall and Kleinhenz 1999, Raffelsiefen 1999). Prosodic phonology assumes different prosodic domains, which are assigned to units as small as the syllable and as large as the whole phonological utterance. The boundary strength of higher-level prosodic domains, e.g. the prosodic phrase, has been shown to affect the acoustic signal. Studies on different phenomena such as, for instance, phrase- and word-final lengthening as well as initial strengthening have illustrated that prosody is predominant in durational variation (see Turk and Shattuck-Hufnagel 2000 for a comprehensive overview), and that investigations of durational variation are insufficient without taking into account the potential impact of higher-ranked prosodic domains. For example, the results in Cambier-Langeveld et al. (1997) and Turk and Shattuck-Hufnagel (2007) reveal that the acoustic duration of word-final syllables and, in certain conditions, main-stressed syllables is lengthened the closer the material is located to a prosodic phrase boundary. Similarly, Fougerson and Keating (1997) found lengthened durations in phrase-initial position, and Cho (2005) reports a stronger articulation of linguistic material at the beginning of prosodic phrases. In the case of NNN compounds, which form one word out of three lexemes, it can be assumed that the material closest to prosodic phrase boundaries, i.e. N1 in phrase-initial position and N3 in phrase-final position, is affected in its acoustic duration.

Morphological factors The presence or absence of morphological boundaries has been shown to affect the acoustic duration of, for instance, words, morphemes and segments. The strongest type of morphological boundary is a word boundary, i.e. the boundary between two separate lexical units, while weaker morphological boundaries are assumed if the involved morphological units form a compound, and if one of the units is an affix. Even different affixes have different degrees of morphological boundary strength, as reported by, for instance, Hay and Plag (2004) and Plag and Baayen (2009). For instance, the suffix *-ee* as in *employee* is assumed to have a lower morphological boundary strength than the suffix *-hood* as in *neighborhood*. Hay and Plag (2004) and Plag and Baayen (2009) demonstrate that this cline of suffix boundary strengths affects the range of attested suffix combinations in the case of multiple affixation, thus forming an acyclic directed graph: While suffixes with higher morphological boundary strengths can be attached to bases that contain a suffix with lower

morphological boundary strength (e.g. in *employeehood*), the reverse case is rarely, if ever, attested.

A similar restriction on attested constituent combinations has been suggested in Baayen (2010) for English compounds. In correspondence to Hay and Plag (2004)'s results that suffixes may be combined only in a particular order, Baayen (2010) finds an equivalent constraint on the combinability of compound constituents. Apparently, there is a preferred order, meaning that a particular constituent may easily occur after another constituent, while the reversed order of the two constituents is dispreferred. If the constraints on combinability of suffixes can be interpreted as differences in morphological boundary strength (Hay and Plag 2004), and if compounds show similar constraints on the combinability of constituents (Baayen 2010), it follows that compounds may also show varying degrees of morphological boundary strength, similar to affixed words.

This conclusion is in line with the argument presented in Kunter and Plag (2016) who claim that in a morphologically complex structure [X Y] Z, where X and Y form an embedded morphological unit, the morphological boundary between X and Y is weaker than the morphological boundary between [XY] and Z. This would also apply to NNN compounds. For instance, in a left-branching compound [N1 N2] N3, the morphological boundary between the embedded constituents N1 and N2 is assumed to be weaker than the morphological boundary between the embedded constituents and the free constituent N3. Simultaneously, a right-branching compound N1 [N2 N3] is considered to have a weaker morphological boundary between the embedded constituents N2 and N3, and a stronger morphological boundary between the free constituent N1 and [N2 N3].

In fact, Lehiste (1972) finds different durations of the stem depending on whether the stem is conjoined with a suffix (e.g. *shaking*) or another free lexeme to form a compound (e.g. *shake-down*). Similarly, Sproat and Fujimura (1993) examined the acoustic properties of the phoneme combination /il/ across different types of boundaries, ranging from non-morphemic boundaries over suffix boundaries to compound boundaries, and eventually word boundaries. They show that different boundary strengths lead to gradient realizations of the segments in focus. In particular, with increasing boundary strength, the degree of velarization of /l/ also increases, while the duration of the rime decreases.

Building on the observation that different degrees of morphological boundary strengths lead to different acoustic realizations, Kunter and Plag (2016) test the prediction that in triconstituent compounds with more than two constituents, embedded forms with weaker morphological boundaries show more phonetic reduction than forms at higher derivational levels. Their experimental data reveal evidence in favor of this hypothesis, as the duration of constituents in their left-branching and right-branching compounds differed even if other variables were accounted for. Yet, this effect interacted with the bigram frequencies of neighboring constituents in the way that high N1N2 bigram frequencies frequently co-occurred in left-branching NNN compounds where N1 and N2 form the embedded compound, while high N2N3 bigram frequencies were mostly found in right-branching NNN compounds with N2 and N3 forming the embedded compound. The origin of the interaction of bigram frequencies

with the branching direction of NNN compounds could not be resolved in Kunter and Plag (2016), possibly due to the rather unbalanced distribution of bigram frequencies in the data set. One possible explanation is that high bigram frequencies in an NNN compound allow for a parsing of the two involved compound constituents as one semantic unit with a closer relation than the one between this semantic unit and the third constituent. This means that the interpretation of semantic relations between the three constituents on the basis of bigram frequencies can be, but need not be congruent with the morphological structure that is suggested by the branching direction.

A smaller, yet not to be underestimated number of NNN compounds shows bigram frequency combinations that work against the branching direction of the compounds: In left-branching NNN compounds, the bigram frequency of the embedded compound N1N2 can be lower than the bigram frequency of N2 and N3 just because the semantics of the compound opposes the lexical frequency of the involved bigrams. Likewise, a right-branching NNN compound can show higher bigram frequencies of N1 and N2 than of the embedded compound N2N3. For these bigram frequency combinations, the constituents show durational patterns that contradict the Embedded Reduction Hypothesis.

This overview implies that the acoustic duration of constituents in NNN compounds is affected by a range of factors. Table 1 summarizes the predictions that emerge from the existing body of research with regard to NNN compounds and their constituents. The analysis presented in the following will test whether these predictions find support in spoken data, which, to our knowledge, has never been done before.

3 Methodology

3.1 Data

In order to test which factors affect constituent durations, we used a data set compiled by Kösling and Plag (2009) for an analysis of prominence assignment in English NNN compounds. The tokens used in that study were extracted from the Boston University Radio Speech Corpus (BURSC, Ostendorf et al. 1997), a collection of speech recordings of professional radio news speakers (3 female and 4 male speakers). The data set includes compounds of both branching directions, i.e. left-branching [N1 N2] N3 and right-branching N1 [N2 N3]. One advantage of the speech data from the BURSC is that this corpus offers recordings of high sound quality, and the speakers, for the most part, provide error-free and fluent speech, which facilitates the segmentation of the recordings. This data set covers a wide range of different bigram frequencies for both branching directions, and is therefore better suited for investigating bigram frequencies and boundary strength as independent factors than the data used in Kunter and Plag (2016) (see section 2).

All NNN compounds meet the following criteria: none of the three constituents is a word from a language other than English (this excludes, for instance, *Hillside hacienda*); contains a possessive clitic (for instance, *tenant's right crisis*); is an initialism (for instance, *U.S. district judge*), or is a proper name (for instance, *Thomas Crown affair*). All compounds

Factors	Predictions
Extra-linguistic	
Speech Rate	The lower the speech rate, the longer the constituent duration. The higher the speech rate, the shorter the constituent duration.
Lexical	
Constituent Frequency	Low constituent frequencies lead to long constituent durations. High constituent frequencies lead to short constituent durations.
Bigram Frequency	Low bigram frequencies of neighboring constituents lead to long constituent durations. High bigram frequencies of neighboring constituents lead to short constituent durations.
Trigram Frequency	Low trigram frequencies lead to long constituent durations. High trigram frequencies lead to short constituent durations.
Phonological Neighborhood Density	A sparse phonological neighborhood density leads to long constituent durations. A dense phonological neighborhood density leads to short constituent durations.
Phonological Levenshtein Distance	A high PLD leads to long constituent durations. A low PLD leads to long constituent durations.
Phonological	
Number of Phonemes	The fewer phonemes, the shorter the constituent duration. The more phonemes, the longer the constituent duration.
Number of Syllables	The fewer syllables, the shorter the constituent duration. The more syllables, the longer the constituent duration.
Accentuation	Accented constituents have longer durations than unaccented constituents.
Prosodic Phrase Boundaries	Following prosodic phrase boundaries lead to longer N3 durations. Preceding prosodic phrase boundaries lead to longer N1 durations.
Morphological	
Morphological Boundaries	Weak morphological boundaries lead to short constituent durations. Strong morphological boundaries do not affect constituent durations.

Table 1: Predicted effects for the acoustic duration of constituents in NNN compounds

are unique in the data set: as discussed in Kösling and Plag (2009), only the first token of any NNN type was sampled, thus avoiding a situation in which particularly frequent NNN types might skew the data set. The data set compiled in this way was made available by Kösling and Plag (2009) for the purpose of the present analysis, together with a classification of the branching direction of each token as well as an acoustic annotation of the compound boundaries and the compound-internal constituent boundaries. The branching direction of about 20 percent of the compounds was determined by Kösling and Plag (2009) on the basis of the orthography: if two constituents were spelled either as one word or were conjoined by a hyphen, these constituents were assumed to form the embedded constituent of the NNN compound. For example, in the NNN compound *weekend series*, the one-word spelling of the first constituents *week* and *end* was considered by the authors to indicate a left-branching compound [*weekend*] *series*, while the hyphen between *whistle* and *blower* was considered a signal of right-branching in *company* [*whistle-blower*]. The branching direction of the remaining 80 percent of tokens was determined by means of a semantic analysis by Kösling and Plag. Only those compounds were included in the data set where both authors unambiguously agreed in their assessment whether the interpretation of the compound suggested left- or right-branching. All in all, the present analysis used 471 NNN compounds from BURSC.

The validity of the acoustic annotations by Kösling and Plag (2009) was tested by the present authors by selecting a random subset of the whole data (about 10%) using the speech analysis software Praat (Boersma and Weenink 2017). For these tokens, the beginning and the end of the compound as well as the two compound-internal constituent boundaries were identified using the common segmentation criteria (cf. Ladefoged and Johnson 2011, Ladefoged and Ferrari Disner 2012). Only the annotation of the compound-final boundary, i.e. the right boundary of N3, showed a significant difference to the annotations provided by Kösling and Plag. Thus, this boundary was re-annotated for the complete data set. All remaining boundaries (i.e. the compound-initial boundary as well as the boundaries between N1 and N2, and between N2 and N3), were retained from Kösling and Plag.

The random sample revealed an inconsistency in the original annotation of identical consonants at constituent boundaries, as, for instance, *Massachusetts state house*, where N1 ends in /s/ and N2 begins with /s/. We set the boundary between the first and the second sound in a doublet (85 items in total) on the basis of the acoustic signal in the spectrogram and the waveform. For instance, a decrease in the activated energy in the spectrogram between two /s/ consonants indicates the end of the previous and the beginning of the following fricative. In cases where this was not possible, we decided to set the boundary at the exact midpoint of the sound doublet.

We excluded 78 constituents from the data (5.5% of a total of 1413 constituents from 471 NNN compounds). Reasons for their exclusion are missing lexical information of constituents (e.g. their phonological neighborhood or PLD), missing pitch measurements that Praat failed to calculate, or extraordinarily high pitch ranges (larger than 20) which we identified as outliers. The final data set comprises 1335 constituents from 465 NNN compounds. The exact distribution of compound constituents in terms of branching direction and speakers is

	F1	F2	F3	M1	M2	M3	M4	TOTAL
left-branching	93	112	230	115	137	45	219	951
right-branching	45	48	96	41	89	6	59	384
TOTAL	138	160	326	156	226	51	278	1335

Table 2: Number of constituents by branching direction and speaker (female speakers: F1–F3, male speakers: M1–M4)

illustrated in Table 2.

3.2 Dependent Variable and Predictor Variables

We used a multifactorial analysis to test the predicted effects (table 1) on the acoustic duration of 1335 compound constituents.

Thus, the dependent variable in our analysis is the ACOUSTIC DURATION in seconds of NNN compounds. We measure durations as the time span between the starting point and the end point of each constituent. The following predictors were included in the statistical model.

Extra-linguistic predictor One speaker-related variable is the overall SPEECHRATE, a measure that expresses the pace at which a speaker produces speech. This was calculated by obtaining the total number of words in the recording containing the NNN compound, subtracting the number of words of the compound itself (which was 2 or 3, depending on the orthography of the compound), and dividing this number by the total duration of the recording minus the duration of the NNN compound. Doing so, item-specific peculiarities, such as an abnormal production of the compound as compared to the rest of the recording, are excluded from this measure.

Lexical predictors Our analysis included three types of lexical frequency: the token frequency of constituents (CONSTITUENTFREQ), the token frequency of two neighboring constituents in the NNN compound (N1N2FREQ and N2N3FREQ), and the token frequency of the NNN compound as a whole (TRIGRAMFREQ). In order to determine the frequency measures, the DVD version of COCA (1990–2012, Davies 08) was queried using the corpus query tool Coquery (Kunter 2017).

We included two measures that indicate the phonological neighborhood of constituents in the analysis in order to account for the organization of the mental lexicon. The phonological neighborhood density (PHONON) and the phonological Levenshtein distance (PLD) of each constituent was collected from the *English Lexicon Project* (henceforth ELP, Balota et al. 2007).

On the basis of previous research investigating the effect of predictability on the acoustic signal, one could assume that the lexical probability of three lexemes occurring together affects acoustic duration. In the case of NNN compounds, however, this type of analysis is

not straightforward as the environment of each compound constituent differs. Testing the lexical probability of constituents in such morphologically complex structures means to look at the probability of a determiner (or, possibly, an adjective) to be followed by N1, N1 to be followed by N2, N2 to be followed by N3, and N3 to be followed by a verb or a prepositional attribute, whatever comes next. This will inevitably lead to very unbalanced probabilities: the probability of one noun out of the class of all nouns in the English lexicon following a determiner like, for instance, 'the' is much smaller than the probability of one certain noun out of the class of all nouns in the English lexicon following another certain noun. Probabilities, thus, appear to be an unreasonable measure for our purposes, as we assume that they cannot capture the relation of lexemes in NNN compounds as n-gram frequencies can.

Phonological predictors The number of phonemes (N_{PHON}) and syllables (N_{SYLL}) for each constituent were also gathered from the ELP (Balota et al. 2007).

In order to account for the presence or absence of pitch accents, we tested different types of pitch-related measures which may be used to incorporate accentuation in a statistical analysis (cf. Kunter 2011, ch. 5 for a discussion). Pitch measures were obtained from the speech recordings using Praat (Boersma and Weenink 2017) with floor and ceiling settings adjusted for the gender of the speakers (female speakers: 100 Hz and 500 Hz; male speakers: 75 Hz and 300 Hz, respectively). They are expressed in semitones (ST) relative to 100 Hz (e.g. a pitch measurement of 100 Hz corresponds to 0 ST, 50 Hz to -12 ST, and 200 Hz to +12 ST). Preliminary modeling revealed that PITCHRANGE, which was calculated as the difference between the highest and the lowest point of the pitch contour of a compound constituent, appears to be the most suitable variable to account for an effect of accentuation on acoustic duration.

All NNN compounds were coded for their phrasal position in the prosodic context. With regard to the preceding phrase boundary PRECBOUNDARY, an NNN compound was labeled 'initial' if it occurred in clause-initial position, otherwise it was labeled 'medial'. Additionally, we coded for a preceding pause by adding '<pause>' to the label, and added '<no pause>' in case of a phrase boundary without an audible pause. Similarly, for the following phrase boundary FOLLBOUNDARY an NNN compound was labeled 'medial' if more linguistic material followed which belonged to the same phrase, otherwise it was labeled 'final'. We added '<pause>' to the label if there was an audible following pause present, while '<no pause>' indicates the absence of a following pause. In our data set NNN compounds with the labels 'initial' and 'final' were always accompanied by a pause. Therefore, only the six combinations shown in table 3 were used as classification in our data set. The rightmost three columns of the table indicate the distribution of the total number of constituents on the types of boundary as well as the distribution of constituents on the types of boundary by branching direction.

Morphological predictors The variable MEMBER corresponds to the constituents in the compound, and has therefore the three levels N1, N2, and N3.

PRECBOUNDARY			N	
<no pause> medial	... at [NP the <i>Boston University school of public health</i>].	1284
<pause> medial	... [NP Those <pause> <i>litmus test supporters</i>] backing Weld...	18
<pause> initial	[NP <pause> <i>Massachusetts school children</i>] are improving ...	33
FOLLBOUNDARY			N	
medial <no pause>	... for [NP immediate <i>emergency state aid for hospitals</i>] .	744
medial <pause>	... preoccupied with [NP the news headlines <pause>] than with ...	95
final <pause>	... in charge of [NP the <i>Mattapan Roxbury area</i> <pause>].	496

Table 3: Prosodic boundaries: three types of boundary with number of tokens (light and dark bars indicate proportion of left-branching and right-branching tokens, respectively)

The morphological structure of the compound is coded in the variable **BRANCHING**. In the analysis, this predictor provides information on the type of NNN compound with the levels *left* and *right*. The branching direction of an NNN compound is naturally related to the variable **MEMBER**. In left-branching compounds, constituents N1 and N2 constitute the embedded compound, and N3 is the free constituent, whereas in right-branching compounds, N1 is the free constituent, and N2 and N3 form the embedded compound (see (1)).

Interactions An interaction of **BRANCHING** $\times$ **MEMBER** in the statistical model reveals whether the effect of **BRANCHING** on the acoustic duration of NNN constituents is the same for all compound constituents N1, N2 and N3, or whether the three constituents of an NNN compound are affected differently. The interaction is therefore inevitable in the analysis of morphological structure and its impact on constituent duration in the statistical model.

In addition, Kunter and Plag (2016) have shown in their analysis of NNN compounds that bigram frequency affected the acoustic duration of embedded constituents and free constituents differently. In order to account for these potential differences, two three-way interactions (**N1N2FREQ** $\times$ **BRANCHING** $\times$ **MEMBER** and **N2N3FREQ** $\times$ **BRANCHING** $\times$ **MEMBER**) were incorporated into our analysis. Doing so, we can investigate the effect of bigram frequencies on embedded constituents (in cases where the bigram corresponds to the embedded constituents in the NNN compound) and on free constituents (in cases where the bigram corresponds to the free constituent and its embedded neighbor in the NNN compound) at the same time.

With regard to the prosodic structure of NNN compounds, we may also assume that the type of prosodic boundary affects embedded constituents and free constituents differently. Therefore, we incorporated the two-way interactions **PRECBOUNDARY** $\times$ **MEMBER** and **FOLLBOUNDARY** $\times$ **MEMBER**. While embedded constituents form one complex unit, free constituents stand on their own. If we consider a prosodic phrase boundary preceding the NNN compound, i.e. which is to the left of the complex word, the branching direction allows us to investigate whether both embedded constituents of a left-branching compound are as affected as the free constituent of a right-branching compound by the prosodic boundary. Simultaneously, we may expect that a prosodic phrase boundary that follows the NNN com-

pound has an effect on the embedded constituents of a right-branching compound, and on the free constituent of a left-branching compound. In order to account for this complex interplay, we could include three-way interactions which are capable to capture these effects. Since the data set does not contain observations for each of the interaction terms, only the three-way interaction `FOLLBORDER` $\times$ `BRANCHING` $\times$ `MEMBER` was included in the model.

Table 4 summarizes all variables that were coded for the purpose of the present analysis, as well as the interactions that were included. Note that each three-way interaction also entails all corresponding lower-order two-way interactions. The table also introduces the two random intercepts `SPEAKER` and `CONSTITUENT`. These random effects reduce the variance introduced by speaker-idiosyncratic patterns or by the random composition of the items.

3.3 Statistical Analysis

The multifactorial analysis was run with the statistical analysis software R (R Core Team 2017). We fitted linear mixed-effects regression models with the `lme4` (Bates et al. 2017) and `lmerTest` (Kuznetsova et al. 2016) packages.

The list of factors in the previous section (table 4) includes two closely related phonological predictors, namely the number of phonemes (`NPHON`) and the number of syllables (`NSYLL`) in a constituent. As both variables express the phonological length of a constituent, they will be strongly correlated, which may introduce undesirable side effects in an analysis. This phenomenon known as collinearity may cause genuine problems in the interpretation of effects of the correlated variables (cf. Baayen 2008, ch. 6 for a discussion). A correlation test confirmed our expectations and revealed a strong monotonic relationship between the two predictors (Spearman’s $\rho = 0.88$) which makes an inclusion of both variables in the regression model unreasonable.

To address this issue, we used a Principal Components Analysis (PCA) that allows us to eliminate potentially harmful effects of this correlation, while at the same time retaining the information contained in `NPHON` and `NSYLL`. As described, for instance, in James et al. (2013, ch. 10.2), a PCA is a statistical transformation that rotates the input variables in such a way that the output variables (the ‘principal components’) minimize the variation in the data space and, used in a regression model, feed in the relevant information of the input variables. The PCA yields as many principal components as there are input variables, but each lower-ranked principal component contributes less to the variation in the data space than any higher-ranked principal component. Table 5 illustrates that `PC1PHON` alone explains 94.3% of the variance in our data set and is as such already a strong predictor in the analysis. `PC2PHON` just exceeds the threshold with a proportion of 5.7% and is therefore also considered a relevant predictor.

The contribution of `NPHON` and `NSYLL` to `PC1PHON` and `PC2PHON` is presented in table 6. Each row shows the loadings of one input variable for each component. Higher absolute values indicate a stronger contribution of the input variable to the principal component, while the sign shows the direction of the contribution. Both phonological variables add identical values to the principal components, although in `PC2PHON` the contribution of `NPHON` goes

Dependent Variable	Range	Median
DURATION	0.089 s to 0.952 s	0.376
Numeric predictors		Median
SPEECHRATE	2.663 wps to 4.235 wps	3 127wps
CONSTITUENTFREQ	207 to 695 121	37 360
N1N2FREQ	0 to 41 150	193
N2N3FREQ	0 to 35 683	35
TRIGRAMFREQ	0 to 29 691	0
PHONOLOGICALNEIGHBORS	0 to 58	5
PLD	1 to 7.100	1.800
NPHON	2 to 14	5
NSYLL	1 to 6	2
PITCHRANGE	0.054 ST to 19.437 ST	5.735
Categorical predictors		
FOLLBOUNDARY	medial <no pause>, medial <pause>, final <pause>	
PRECBOUNDARY	<no pause> medial, <pause> medial, <pause> initial	
MEMBER	N1, N2, N3	
BRANCHING	left, right	
Interactions		
FOLLBOUNDARY × BRANCHING × MEMBER		
N1N2FREQ × BRANCHING × MEMBER		
N2N3FREQ × BRANCHING × MEMBER		
Random Effects		
SPEAKER		
CONSTITUENT		

Table 4: Variables in the initial regression model

Importance of components:		
	PC1 _{PHON}	PC2 _{PHON}
Standard deviation	1.373	0.338
Proportion of Variance	0.943	0.057
Cumulative Proportion	0.943	1.000

Table 5: Principal component analysis for phonological measurements

	PC1 _{PHON}	PC2 _{PHON}
N _{PHON}	-0.707	0.707
N _{SYLL}	-0.707	-0.707

Table 6: Loadings of principal components by phonological variables

in the positive direction while the contribution of N_{SYLL} is negative.

Another issue in linear models concerns the assumption that the residuals are normally distributed. However, visual inspection of the residuals in the initial model revealed a non-normal distribution, which was confirmed by a Kolmogorov-Smirnov test. In order to alleviate this deviation, the dependent variable DURATION was Box-Cox-transformed using an exponent of $\lambda = 0.59$ (Box and Cox 1964). This transformation removed the non-normality of the residuals.

The model was simplified by stepwise exclusion of insignificant predictors (cf. Baayen 2008, ch. 7). For each predictor that did not reach the significance level of $p \leq 0.05$, and which was not part of a higher-order interaction, a new model was fitted that excluded this predictor. If a pairwise ANOVA did not reveal a significantly worse fit of the new model, that predictor was eliminated from the analysis. This was the case for PHONOLOGICALNEIGHBORS, PLD, PRECBOUNDARY, and TRIGRAMFREQ. The final model included all predictors listed in table 7.

For the interpretation of interaction effects, we used the `diffsmeans` function from the `lmerTest` package, which produces contrasts of categorical variables in a linear mixed-effects model. The function revealed which interaction terms involving factorial predictors in our statistical model show significant contrasts. This is a helpful tool for the interpretation of statistical effects and their directions.

4 Results

The reference level for the Intercept of the regression model is PRECBOUNDARY = <no pause> medial, FOLLBOUNDARY = medial <no pause>, BRANCHING = left and MEMBER = N1. The coefficients indicate the estimated effect of the respective predictor on constituent duration. Positive numbers show that the variable in question increases the constituent duration, and negative numbers indicate a decrease. Note that all coefficients relate to the Box-Cox-transformed durations. For the convenience of the reader, we back-transformed the dependent variable to duration in seconds in figures 1 and 2 which illustrate

the partial effects of the predictors.

Single effects Figure 1 displays the main effects of `CONSTITUENTFREQ`, `PITCHRANGE`, `SPEECHRATE` and `PRECBOUNDARY` with back-transformed durations in seconds on the y-axis. In each plot, only the predictor under investigation is varied, while all other predictors are held constant at their median (for numeric predictors). The top left plot for `CONSTITUENTFREQ` reveals the expected effect: Higher constituent frequencies lead to shorter durations.

The effect of `PITCHRANGE` is shown in the top right plot. Pitch ranges are indicated on the x-axis as the difference of minimum pitch and maximum pitch within one constituent in semitones. As pitch range is an operationalization of accentuation, the plot indicates that accented constituents with larger pitch ranges have longer durations than unaccented constituents with smaller pitch ranges.

The bottom left plot illustrates the effect of `SPEECHRATE` on the x-axis. As can be seen, constituent durations decrease with increasing speech rate, i.e. with a faster pace of a speaker to produce speech the constituents get shorter.

In the bottom right plot the effect of `PRECBOUNDARY` on acoustic duration is shown. The x-axis indicates the three categorical levels ‘clause-initial <pause>’, ‘phrase-medial <no pause>’ and ‘phrase-medial <pause>’ from left to right. The prosodically strongest boundary ‘clause-initial <pause>’ on the left shows the smallest effect on constituent durations, while constituents are longest at a ‘phrase-medial <pause>’ boundary.

The two plots in figure 2 illustrate the effects of the two phonological principal components `PC1PHON` and `PC2PHON` with constituent durations in seconds on the y-axis. The left plot shows that with increasing `PC1PHON` the constituent durations are shortened. Similarly, the right plot indicates that an increase of the second principal component `PC2PHON` leads to a decrease of constituent durations.

The negative correlation of the principal components with the input variables `NPHON` and `NSYLL` explains the direction of these effects: higher PC values indicate a smaller number of phonemes and syllables in the constituent. Therefore, the decreasing regression line reveals longer durations for constituents with a relatively high number of phonological units, and shorter durations for constituents with fewer phonemes and syllables, respectively. Crucially, the effect of `PC1PHON` is the strongest of all main effects. This can be seen in the steeply decreasing regression line in the plot for `PC1PHON` which covers the whole range of constituent durations on the y-axis: Constituents with minimal values for `PC1PHON` have durations of approx. 700 ms, while constituents with maximal `PC1PHON` values are approx. 200 ms long. In contrast, the smaller effect sizes of the predictors `CONSTITUENT FREQ`, `PITCHRANGE`, `SPEECHRATE` and `PC2PHON` on constituent durations range between approx. 80 ms and 120 ms.

FOLLBOUNDARY × BRANCHING × MEMBER The regression analysis reveals that the three-way interaction `FOLLBOUNDARY × BRANCHING × MEMBER` has a significant effect on constituent durations. The effect of the interaction is illustrated in figure 3. The top plot

	Estimate	Std. Error	d.f.	t	p
(Intercept)	0.7056	0.0288	951.1	24.489	< 0.001 ***
log CONSTITUENTFREQ	-0.0085	0.0017	599.7	-5.045	< 0.001 ***
PC1PHON	-0.0573	0.0017	665.5	-33.019	< 0.001 ***
PC2PHON	-0.0329	0.0061	580.7	-5.388	< 0.001 ***
PITCHRANGE	0.0055	0.0005	1217.0	10.323	< 0.001 ***
SPEECHRATE	-0.0279	0.0061	876.3	-4.580	< 0.001 ***
FOLLBOUNDARY = medial <pause>	0.0256	0.0114	901.7	2.236	0.026 *
FOLLBOUNDARY = final <pause>	0.0131	0.0061	1017.5	2.147	0.032 *
PRECBOUNDARY = medial <no pause>	0.0186	0.0093	1048.7	1.988	0.047 *
PRECBOUNDARY = medial <pause>	0.0299	0.0164	1164.9	1.816	0.070 *
log N1N2FREQ	-0.0032	0.0014	1285.3	-2.327	0.020 *
log N2N3FREQ	-0.0024	0.0013	1139.5	-1.786	0.074
BRANCHING = right	-0.0126	0.0169	1093.102	-0.746	0.456
MEMBER = N2	-0.0081	0.0132	1294.1	-0.614	0.539
MEMBER = N3	0.0167	0.0126	1292.6	1.317	0.188
FOLLBOUNDARY = medial <pause> × MEMBER = N2	-0.0268	0.0157	887.1	-1.705	0.089
FOLLBOUNDARY = medial <pause> × MEMBER = N3	-0.0197	0.0168	1185.0	-1.175	0.240
FOLLBOUNDARY = final <pause> × MEMBER = N2	-0.0112	0.0086	1028.4	-1.303	0.193
FOLLBOUNDARY = final <pause> × MEMBER = N3	0.0166	0.0091	1205.6	1.835	0.0667
BRANCHING = right × MEMBER = N2	-0.0072	0.0239	1212.3701	-0.303	0.762
BRANCHING = right × MEMBER = N3	0.0331	0.0229	1164.0151	1.445	0.149
FOLLBOUNDARY = medial <pause> × BRANCHING = right	-0.0442	0.0221	938.1	-1.998	0.046 *
FOLLBOUNDARY = final <pause> × BRANCHING = right	-0.0033	0.0114	945.6	-0.286	0.775
log N1N2FREQ × BRANCHING = right	0.0014	0.0028	1122.2	0.480	0.631
log N1N2FREQ × MEMBER = N2	-0.0041	0.0019	1293.6	-2.096	0.036 *
log N1N2FREQ × MEMBER = N3	0.0045	0.0018	1276.3	2.446	0.015 *
log N2N3FREQ × BRANCHING = right	0.0034	0.0023	982.6	1.489	0.137
log N2N3FREQ × MEMBER = N2	0.0054	0.0019	1168.2	2.818	0.005 **
log N2N3FREQ × MEMBER = N3	-0.0014	0.0019	1244.6	-0.708	0.479
FOLLBOUNDARY = medial <pause> × BRANCHING = right × MEMBER = N2	0.0435	0.0312	937.6	1.393	0.164
FOLLBOUNDARY = medial <pause> × BRANCHING = right × MEMBER = N3	0.0652	0.0311	1047.2	2.095	0.036 *
FOLLBOUNDARY = final <pause> × BRANCHING = right × MEMBER = N2	0.0114	0.0162	1042.2	0.704	0.482
FOLLBOUNDARY = final <pause> × BRANCHING = right × MEMBER = N3	0.0140	0.0165	1120.8	0.846	0.398
log N1N2FREQ × BRANCHING = right × MEMBER = N2	0.0104	0.0039	1172.6	2.688	0.007 **
log N1N2FREQ × BRANCHING = right × MEMBER = N3	-0.0034	0.0038	1181.4	-0.899	0.369
log N2N3FREQ × BRANCHING = right × MEMBER = N2	-0.0071	0.0033	1113.5	-2.149	0.032 *
log N2N3FREQ × BRANCHING = right × MEMBER = N3	-0.0072	0.0033	1172.4	-2.195	0.028 *

Table 7: Fixed effects in final regression model (dependent variable: transformed constituent duration, N = 1335, reference level of the Intercept: FOLLBOUNDARY = NP-medial, BRANCHING = left, MEMBER = N1)

Figure 1: Effects of constituent frequency, pitch range and speech rate (back-transformed, in sec)

Figure 2: Effects of PC1phon and PC2phon (back-transformed, in sec)

displays the effect of $\text{FOLLBOUNDARY} \times \text{MEMBER}$ for left-branching compound constituents, while the plot on the bottom shows the effect of this interaction on constituents in right-branching compounds. The y-axes indicate back-transformed durations in seconds, and the three panels for each prosodic position of the NNN compound indicate the constituents N1, N2 and N3 on the x-axes. In addition to the regression table and the interaction plot, the function `diffsmeans` informs about significant contrasts of interaction terms from the regression model. For each interaction term of $\text{FOLLBOUNDARY} \times \text{BRANCHING} \times \text{MEMBER}$ this function indicates whether the difference of two terms leads to significantly different constituent durations.

In left-branching compounds, N2 constituent durations are shorter than N1 and N3 constituents across all three prosodic boundaries, and with following pauses this effect is even significant. In phrase-final position, N3 constituents are the significantly longest of all three compound constituents. At the same time, N1 constituent durations become significantly longer than N2 constituents with a following pause. This means that the predicted effect of following prosodic boundaries is visible in left-branching compounds, although it is not only compound-final N3 constituents that are affected but also N1 constituents.

The results for right-branching compounds draw a different picture. At different prosodic boundaries, constituent durations become significantly different: While without a following pause, all three compound constituents are of the same length, N3 constituents become significantly longer than N1 constituents in NP-medial position with a following pause, while N2 constituents pattern in between. In phrase-final position, both N1 and N2 constituent durations are significantly shorter than N3 constituents. In right-branching compounds, thus, the effect of following prosodic boundaries on compound-final N3 constituent durations is as expected.

N1N2FREQ $\times$ BRANCHING $\times$ MEMBER and N2N3FREQ $\times$ BRANCHING $\times$ MEMBER Before turning to the three-way interactions involving bigram frequencies, we will first take a glimpse at the constituent durations in an average left- and right-branching compound, i.e. one in which all numeric variables are set to the arithmetic mean and which occurs in the most frequent prosodic position, namely "NP-medial". As shown in figure 4, N1 and N2 constituents in left-branching compounds are shorter than N3 constituents. Similarly, N1 and N2 constituents in right-branching compounds have equal length and are significantly shorter than N3 constituents. The two significant three-way interactions indicate that this average distribution changes with changing bigram frequencies, but these changes affect the individual constituents in different ways. Furthermore, the effect is different depending on the branching direction. Figure 5 illustrates the effect of N1N2FREQ on the three constituents in left- and right-branching compounds, while figure 6 shows the corresponding effect of N2N3FREQ.

Each bar in figure 5 illustrates the average deviation of a constituent from the average duration of that constituent as shown in figure 4 (N1 = blue, N2 = red, N3 = green). N1N2FREQ is equivalent to the frequency of the embedded compound N1N2 in left-branching

Figure 3: Effect of different positions in a phrase on constituent durations (back-transformed, in sec), top: left-branching, bottom: right-branching

Figure 4: Constituent duration in left- and right-branching compounds (back-transformed, in sec)

Figure 5: Effect of low, medium and high N1N2 frequencies on the deviation of constituent durations from an average constituent (back-transformed, in sec)

compounds, while it is the bigram frequency of the free constituent N1 and the neighboring embedded constituent N2 in right-branching compounds. There are three panels in the plot indicating three levels of N1N2 bigram frequencies: low, medium and high N1N2 bigram frequencies, indicating the contribution of N1N2FREQ at the 25% quantile, the median, and the 75% quantile. The branching direction is indicated in the bars by lighter (left-branching) and darker (right-branching) shades of the respective colors. The zero line on the y-axis indicates the average duration of constituents from which two colored bars per constituent depart. In order to see the full picture of the effect of the three-way interaction involving N1N2FREQ, the deviation from an average constituent in figure 5 is added or subtracted from the constituent duration indicated in figure 4, depending on the direction of the deviation.

The strongest effect of N1N2 bigram frequencies on constituent duration can be seen in the deviations of N2 from the duration of an average N2 constituent. In left-branching compounds, N2 significantly decreases in duration with higher N1N2 bigram frequencies, which is illustrated by the light-red bars spreading in the negative direction: with high N1N2FREQ, left-branching N2 constituents are shorter than an average left-branching N2 constituent (approx. -60 ms). At the same time, N2 durations increase significantly in right-branching compounds with increasing N1N2 bigram frequencies, as indicated by the dark-red bars going in the positive direction: with high N1N2FREQ, N2 constituents are longer than an average right-branching N2 (approx. +45 ms).

The duration of left-branching N1 constituents is also significantly affected by an increase of N1N2 bigram frequency (approx. -25 ms), whereas N3 constituents are hardly affected. Likewise, right-branching N1 and N3 constituent durations do not change with increasing N1N2 bigram frequencies. Overall, increasing N1N2 bigram frequencies produce more striking duration differences in left-branching compound constituents than in right-branching

Figure 6: Effect of low, medium and high N2N3 frequencies on the deviation of constituent durations from an average constituent (back-transformed, in sec)

compound constituents.

Figure 6 illustrates the effect of N2N3 bigram frequencies and branching direction on constituent durations. Similar to figure 5, the zero line marks the average duration of constituents, and it is the point of departure for the colored bars indicating the three constituents N1 (blue), N2 (red) and N3 (green) in two different color shades indicating left-branching (light) and right-branching (dark) constituents. The three panels show three levels of N2N3FREQ, namely low, medium and high N2N3 bigram frequencies, while N1N2 bigram frequencies are kept constant at $\log(N1N2frequency) = 0$ in the intercept of the regression model. N2N3FREQ is equivalent to the frequency of the embedded compound in right-branching compounds, while it is the bigram frequency of the embedded constituent N2 and the free constituent N3 in left-branching compounds.

The most striking effect of this interaction can be seen in N3 constituent durations. With increasing N2N3FREQ, right-branching N3 constituents decrease significantly in duration compared to an average right-branching N3 constituent (approx. -60 ms). The effect on left-branching N3 constituents is similar, however, these constituents are not as affected as their right-branching counterparts (approx. -30 ms). With regard to N1 and N2 durations, there is a significant effect of N2N3 bigram frequencies on left-branching N2 constituents: Increasing N2N3FREQ lead to longer N2 durations (approx. +25 ms), whereas left-branching N1 constituents as well as right-branching N1 and N2 constituents are hardly affected.

5 Discussion and Conclusion

The results of our statistical analysis provide support for the majority of predictions from section 3 (cf. table 1). However, there are also unexpected effects and interactions we will

discuss in the following.

SPEECHRATE affects constituent durations in the expected manner: a high speech rate, i.e. fast speech, leads to shorter constituent durations, while a low speech rate results in longer constituent durations. A higher number of phonological units, encoded in PC1PHON and PC2PHON, leads to longer constituent durations, which was predicted. These findings are not surprising, yet the fact that they correspond to the results of previous studies is promising as it suggests that our data set is well-suited for an investigation of more complex factors influencing acoustic duration. The impact of PC1PHON is particularly intriguing, since it has by far the strongest effect on acoustic duration that is observable in the steeply falling regression line in figure 2 that covers duration differences of up to approx. 500 ms.

Furthermore, the operationalization of prominence in the form of PITCHRANGE proved to be successful. Our analysis reveals that this variable correlates with acoustic duration in the sense that more prominent compound constituents, i. e. those with a larger PITCHRANGE, have a longer constituent duration. This confirms what previous studies have found with regard to the impact of pitch accents on acoustic duration (cf. Turk and Sawusch 1996; Turk and White 1999; de Jong 2004; Kunter 2011; Morrill 2012). We can assume, thus, that PITCHRANGE is a reasonable choice for a predictor variable which is supposed to represent prominence patterns in NNN compounds.

With regard to prosodic phrase boundaries we can state that NNN compounds behave differently according to their morphological structure, although not necessarily in the expected manner. Our results regarding preceding prosodic boundaries, thus, do not provide evidence for an effect of initial strengthening that has been reported in, for instance, Fougeron and Keating (1997). In fact, compound constituents are lengthened if they are preceded by phrase-medial pauses, but they are shortest clause-initially. This finding, however, should be taken with a pinch of salt: With only 51 constituents (3.8% of the data), the total number of pauses preceding the NNN compounds, which is in our study indicative for a stronger prosodic phrase boundary, is very small. It is not impossible that an effect of the preceding boundary type emerges in data sets with a more balanced distribution of this variable.

Our results regarding the influence of following prosodic phrase boundaries suggest that compound constituents of both branching directions are strongly affected by this variable. In right-branching compounds, a following prosodic phrase boundary particularly affects N3 constituents in such a way that they get longer in the presence of a stronger boundary, whereas the effect on N1 and N2 constituents is weaker. With regard to left-branching compounds, it is not only N3 constituents, but also N1 constituent durations that become longer with stronger prosodic boundaries. Overall, these results confirm the predictions we made for prosodic phrase boundaries: Based on previous research (cf. Cambier-Langeveld et al. 1997; Turk and Shattuck-Hufnagel 2007), we predicted that stronger prosodic boundaries yield longer N3 constituent durations. Unexpectedly, there is another significant lengthening effect of prosodic boundaries on N1 durations.

The significance of the three-way interaction FOLLBOUNDARY $\times$ BRANCHING $\times$ MEMBER points at two main findings. First, prosodic domains are capable to affect the phonetic signal

of linguistic material, even if they do so in an unexpected manner. Second, and this is a notable finding with regard to complex words, the branching direction appears to play a role in how prosodic domains act upon acoustic duration. As mentioned before, left-branching compound constituents are hardly affected by the type of prosodic phrase boundary that follows the NNN compound, whereas right-branching N1 constituent durations increase. This robust effect indicates a correlation of prosodic structure with the internal structure of NNN compounds, which is, to our knowledge, a rather unexplored field. It runs counter to some of the traditional assumptions of phonological theory, such as the concept of Bracket Erasure from Lexical Phonology (cf. Kiparsky 1982).

Both PLD and PHONON, which indicate the number of similar and related lexical entries that a compound constituent has in the mental lexicon, do not affect constituent durations. This means that it does not make a difference for acoustic duration whether speakers produce a lexeme with a high number of related entries, or a lexeme which is rather isolated in the mental lexicon. We have been cautious with the prediction for PLD, since only a few previous studies have used this lexical measure and investigated data from perception tasks (cf. Suárez et al. 2011), while we are investigating production data from a speech corpus. Phonological neighborhood density, on the other hand, is considered an established correlate of the phonetic signal (cf. Gahl et al. 2012) in the sense that higher densities lead to more phonetic reduction. Yet, this predictor did not reach statistical significance in our analysis. One possible explanation for this surprising finding is the notable difference in the distribution of phonological neighbors: the median of neighborhood density in Gahl et al. (2012)’s analysis is 21, while in our data set the median is 5.

Similarly, we have found that TRIGRAMFREQ does not affect constituent durations. Based on previous research on n-gram frequency effects, we predicted a shortening effect of the frequency of the three compound constituents to occur together, however, the small range of trigram frequencies appears to prevent such an impact. In our data, trigram frequencies range from 0 to 29691 tokens in COCA, with a median of 0 and a mean of 304.4. These figures suggest that the majority of NNN compounds rarely occurs in the corpus. Therefore, the fact that we do not find a significant effect of trigram frequency on acoustic duration does not contradict findings of, for instance, Arnon and Cohen Priva (2013), who report a shortening effect of n-gram frequencies on the duration of high-frequency phrases (*a lot of work*, mean frequency = 12.04 per million) and low-frequency phrases (*a lot of years*, mean frequency = 2.25 per million). Rather, the potential explanatory power of this lexical variable is weakened by particularly small token frequencies.

Contrary to this, an increasing CONSTITUENTFREQ has a robust shortening effect on acoustic duration. Our results support findings from previous research that has investigated low-frequent and high-frequent linguistic units and their phonetic signal which is presumably shaped by their frequency (cf. Wright 1979, Pluymaekers et al. 2005, Gahl 2008). Inevitably, the question arises what this effect actually tells us about the processing of compounds. After all, NNN compounds are morphologically complex units which contain smaller units, namely free constituents and embedded compounds, which in turn consist of two constituents. Thus,

it seems that speakers accomplish an intricate task when processing NNN compounds: token frequencies of constituents affect their durations even though the constituents are not loosely ordered individual lexemes, but form an NNN compound with a certain internal organization.

The interplay of $N1N2FREQ \times BRANCHING \times MEMBER$ and $N2N3FREQ \times BRANCHING \times MEMBER$ seems to require even more effort on the part of speakers. Originally, we predicted an effect of the morphological structure, i. e. branching direction, on the duration of compound constituents as well as a decrease of constituent durations with increasing bigram frequencies. The robust effect of the three-way interactions, however, indicates that the impact of morphological structure and lexical frequency can only be reasonably investigated when simultaneously taking both factors into account.

In left-branching compounds, where N1 and N2 constituents form the embedded compound, we expected shorter constituent durations for N1 and N2 and comparatively longer constituent durations for the free constituent N3. Likewise, with high N1N2 bigram frequencies, we predicted shorter N1 and N2 constituent durations while N3 constituents remain unaffected. As a consequence, both predicted effects are assumed to manifest themselves in this duration pattern: $N1, N2 < N3$. In fact, this is more or less what we see: Except for N1 constituents, which are hardly affected by N1N2 bigram frequencies, we observe that N1 and N2 constituent durations are overall shorter than N3 constituent durations.

Similarly, for right-branching compounds with N2 and N3 constituents forming the embedded compound and N1 as the free constituent we predicted shorter constituent durations for N2 and N3 and comparatively longer constituent durations for N1. We predicted that high N2N3 bigram frequencies lead to shorter N2 and N3 constituent durations, while they have no effect on N1 constituents. Again, the duration pattern that follows from our predictions is the same for both factors: $N1 > N2, N3$. This, however, is not what we find in our results. Instead, similar to left-branching compounds, N3 constituents are overall longer than N1 and N2 constituents in right-branching compounds. Thus, the Embedded Reduction Hypothesis from Kunter and Plag (2016) makes wrong predictions about the acoustic duration in NNN compounds.

The effect that we observe of N1N2 bigram frequencies on the embedded constituent N2 in right-branching compounds as well as of N2N3 bigram frequencies on the embedded constituent N2 in left-branching compounds is even more complex. Apparently, the token frequency of bigrams that are not the embedded compound of an NNN compound affects constituent durations in an unexpected manner. Increasing N1N2 bigram frequencies lead to N2 lengthening in right-branching compounds, but to N2 shortening in left-branching compounds. An increase of N2N3 bigram frequencies leads to a lengthening of left-branching N2 constituents, whereas N2 in right-branching compounds is not affected at all. This means that the acoustic duration of N2 constituents, which in both cases form part of the embedded compound, is affected by the bigram frequency that does not coincide with the morphological structure of the NNN compound. This is surprising for two reasons. First, the morphological structure of NNN compounds actually does not care about the semantic relation of compound constituents which do not form a morphological unit together. Second, lexical frequency is

usually assumed to have a unidirectional effect on acoustic duration: If linguistic material is more frequent, its acoustic duration is shorter because access to the lexical entry in the mental lexicon is simplified. Yet, our findings show that the impact of morphological structure on the one hand and of lexical frequency on the other hand operate in different ways.

We interpret these unexpected findings as an attempt of speakers to disambiguate the morphological structure of an NNN compound from the frequency-driven parsing of the NNN compound. In cases where the branching direction and a highly frequent bigram do not coincide, speakers seem to modify the duration pattern in such a way that hearers are able to parse the intended branching direction of the NNN compound despite the bigram frequency. The right-branching compound ‘consumer credit chain’, for instance, has a N2N3 bigram of 0 for ‘credit chain’, and a N2N3 bigram frequency of 247 for ‘consumer credit’. Yet, the compound does not refer to a consumer credit having to do with chains, but rather with a credit chain having to do with consumers. Thus, the lengthening of N2 constituents in right-branching compounds when N1N2 bigram frequencies are high is a speaker’s attempt to clarify that the compound is not left-branching with short N1N2 constituents, but that it is in fact right-branching.

The preceding analysis has confirmed that the acoustic duration of words is affected by a variety of variables, such as the presence of a pitch accent, the number of segments and syllables, the speaking rate at which speakers produce words, and the frequency of a constituent. Most of these variables have been found to contribute to word duration in previous research that did not explicitly look at morphologically complex words, and our findings for NNN compounds match these results. More intriguingly, however, are the various interactions of variables that reveal a notable influence of the internal structure of the NNN compounds. The question remains whether the parsing of complex structures such as NNN compounds is primarily frequency-driven or morphology-driven, or whether these factors operate in parallel. It seems that an effect of lexical n-gram frequency is more general in the sense that the n-gram involved is not variable, while an effect of branching direction is: An n-gram can take diverse positions and, thus, hold different roles in a complex word. This means that, for instance, one and the same bigram can constitute the embedded compound in an NNN compound as well as hold the positions of the free constituent and the neighboring embedded constituent, resulting in two distinct roles in the NNN compound. Besides, the semantics of NNN compounds is determined by the internal relations of compound constituents, but which type of relations is superior: the morphological relations involving different degrees of boundary strength and rules according to which constituents can be ordered, or the lexical relations determined by the frequency of occurrence of constituents? At this point, we cannot provide an answer to this question, and it remains still unclear whether there is in fact a clear-cut answer.

References

- Arnon, I. and Cohen Priva, U. (2013). More than words: The effect of multi-word frequency and constituency on phonetic duration. *Language and speech*, 56(Pt 3):349–371.
- Aylett, M. and Turk, A. (2006). Language redundancy predicts syllabic duration and the spectral characteristics of vocalic syllable nuclei. *Journal of the Acoustical Society of America*, 119(5):3048–3058.
- Baayen, H. R. (2008). *Analyzing linguistic data: A practical introduction to statistics using R*. Cambridge Univ. Press, Cambridge.
- Baayen, R. H. (2010). The directed compound graph of English: An exploration of lexical connectivity and its processing consequences. In Olsen, S., editor, *New Impulses in Word-Formation*, number 17 in *Linguistische Berichte*, page 383–402. Buske, Hamburg.
- Balota, D. A., Yap, M. J., Cortese, M. J., Hutchison, K. A., Kessler, B., Loftis, B., Neely, J. H., Nelson, D. L., Simpson, G. B., and Treiman, R. (2007). The English Lexicon Project. *Behavior Research Methods*, 39(3):445–459.
- Bates, D., Mächler, M., Bolker, B., and Walker, S. (2017). lme4: Linear mixed-effects models using Eigen and S4.
- Bauer, L. (2017). *Compounds and Compounding*. Cambridge Studies in Linguistics. Cambridge University Press.
- Bell, A., Brenier, J. M., Gregory, M., Girand, C., and Jurafsky, D. (2009). Predictability effects on durations of content and function words in conversational English. *Journal of Memory and Language*, 60(1):92–111.
- Bell, A., Jurafsky, D., Fosler-Lussier, E., Girand, C., Gregory, M., and Gildea, D. (2003). Effects of disfluencies, predictability, and utterance position on word form variation in English conversation. *The Journal of the Acoustical Society of America*, 113(2):1001–1024.
- Boersma, P. and Weenink, D. (2017). Praat: Doing phonetics by computer.
- Box, G. and Cox, D. R. (1964). An analysis of transformations. *Journal of the Royal Statistical Society*, 26(2):211–256.
- Cambier-Langeveld, G. M., Nespor, M. A., and van Heuven, V. J. (September 22-25, 1997). The domain of final lengthening in production and perception in Dutch.
- Cho, T. (2005). Prosodic strengthening and featural enhancement: Evidence from acoustic and articulatory realizations of /a,i/ in English. *The Journal of the Acoustical Society of America*, 117(6):3867–3878.
- Davies, M. (2008–). The Corpus of Contemporary American English (COCA): 520 million words, 1990-present.

- de Jong, K. (2004). Stress, lexical focus, and segmental focus in English: Patterns of variation in vowel duration. *Journal of Phonetics*, 32(4):493–516.
- Farnetani, E., Torsello, C. T., and Cosi, P. (1988). English compound versus non-compound noun phrases in discourse: An acoustic and perceptual study. *Language and Speech*, 31(2):157–180.
- Fougeron, C. and Keating, P. A. (1997). Articulatory strengthening at edges of prosodic domains. *The Journal of the Acoustical Society of America*, 101(6):3728–3740.
- Gahl, S. (2008). Time and thyme are not homophones: The effect of lemma frequency on word durations in spontaneous speech. *Language*, 84(3):474–496.
- Gahl, S., Yao, Y., and Johnson, K. (2012). Why reduce? Phonological neighborhood density and phonetic reduction in spontaneous speech. *Journal of Memory and Language*, 66(4):789–806.
- Giegerich, H. J. (2004). Compound or phrase? English noun-plus-noun constructions and the stress criterion. *English Language and Linguistics*, 8(1):1–24.
- Gregory, M., Raymond, W. D., Bell, A., Fosler-Lussier, E., and Jurafsky, D. (1999). The effects of collocational strength and contextual predictability in lexical production. *Chicago Linguistic Society*, 35:151–166.
- Hall, T. A. and Kleinhenz, U., editors (1999). *Studies on the Phonological Word [Conference on the Phonological Word, Berlin, 24.-26.10.1997]*, volume 174 of *Amsterdam studies in the theory and history of linguistic science Series 4, Current issues in linguistic theory*. Benjamins, Amsterdam.
- Hay, J. B. and Plag, I. (2004). What constrains possible suffix combinations? On the interaction of grammatical and processing restrictions in derivational morphology. *Natural Language & Linguistic Theory*, 22(3):565–596.
- James, G., Witten, D., Hastie, T., and Tibshirani, R. (2013). *An introduction to statistical learning with applications in R*. Springer, New York, 6th printing 2015 edition.
- Jurafsky, D., Bell, A., Gregory, M., and Raymond, W. D. (2001). Probabilistic relations between words: Evidence from reduction in lexical production. In Bybee, J. L. and Hopper, P. J., editors, *Frequency and the Emergence of Linguistic Structure*, volume 45 of *Typological Studies in Language*, pages 229–254. John Benjamins Publishing Company, Amsterdam.
- Kiparsky, P. (1982). Lexical morphology and phonology. In Linguistic Society of Korea, editor, *Linguistics in the morning calm*, pages 3–99. Hanshin Pub. Co, Seoul.
- Kösling, K., Kunter, G., Baayen, H., and Plag, I. (2013). Prominence in triconstituent compounds: Pitch contours and linguistic theory. *Language and speech*, 56(Pt 4):529–554.

- Kösling, K. and Plag, I. (2009). Does branching direction determine prominence assignment? An empirical investigation of triconstituent compounds in English. *Corpus Linguistics and Linguistic Theory*, 5(2).
- Kunter, G. (2011). *Compound stress in English: The phonetics and phonology of prosodic prominence*, volume 539 of *Linguistische Arbeiten*. De Gruyter, Berlin.
- Kunter, G. (2017). Coquery: A free corpus query tool.
- Kunter, G. and Plag, I. (2016). Morphological embedding and phonetic reduction: The case of triconstituent compounds. *Morphology*, 26(2):201–227.
- Kuznetsova, A., Brockhoff, P. B., and Bojesen Christensen, R. H. (2016). lmerTest: Tests in linear mixed effects models.
- Ladefoged, P. and Ferrari Disner, S. (2012). *Vowels and consonants*. Blackwell Publishing Ltd., Chichester, 3. edition.
- Ladefoged, P. and Johnson, K. (2011). *A course in phonetics*. Wadsworth, Boston, 6. edition.
- Lehiste, I. (1972). The timing of utterances and linguistic boundaries. *The Journal of the Acoustical Society of America*, 51(6B):2018–2024.
- Liberman, M. and Prince, A. (1977). On stress and linguistic rhythm. *Linguistic Inquiry*, 8(2):249–336.
- Morrill, T. (2012). Acoustic correlates of stress in English adjective-noun compounds. *Language and speech*, 55(Pt 2):167–201.
- Ostendorf, M., Price, P., and Shattuck-Hufnagel, S. (1997). Boston University Radio Speech Corpus: Speech and annotations.
- Plag, I. and Baayen, R. H. (2009). Suffix ordering and morphological processing. *Language*, 85(1):109–152.
- Pluymaekers, M., Ernestus, M., and Baayen, R. H. (2005). Lexical frequency and acoustic reduction in spoken Dutch. *The Journal of the Acoustical Society of America*, 118(4):2561–2569.
- R Core Team (2017). R: A language and environment for statistical computing.
- Raffelsiefen, R. (1999). Diagnostics for prosodic words revisited. In Hall, T. A. and Kleinhenz, U., editors, *Studies on the Phonological Word*, volume 174 of *Current Issues in Linguistic Theory*, page 133. John Benjamins Publishing Company, Amsterdam.
- Selkirk, E. O. (1980). The role of prosodic categories in English word stress. *Linguistic Inquiry*, 11(3):563–605.
- Seyfarth, S. (2014). Word informativity influences acoustic duration: Effects of contextual predictability on lexical representation. *Cognition*, 133(1):140–155.

- Sproat, R. and Fujimura, O. (1993). Allophonic variation in English /l/ and its implications for phonetic implementation. *Journal of Phonetics*, 21:291–311.
- Suárez, L., Tan, S. H., Yap, M. J., and Goh, W. D. (2011). Observing neighborhood effects without neighbors. *Psychonomic Bulletin & Review*, 18(3):605–611.
- Turk, A. E. and Sawusch, J. R. (1996). The processing of duration and intensity cues to prominence. *The Journal of the Acoustical Society of America*, 99(6):3782–3790.
- Turk, A. E. and Shattuck-Hufnagel, S. (2000). Word-boundary-related duration patterns in English. *Journal of Phonetics*, 28(4):397–440.
- Turk, A. E. and Shattuck-Hufnagel, S. (2007). Multiple targets of phrase-final lengthening in American English words. *Journal of Phonetics*, 35(4):445–472.
- Turk, A. E. and White, L. (1999). Structural influences on accentual lengthening in English. *Journal of Phonetics*, 27(2):171–206.
- Uhmann, S. (1992). Contextualizing relevance: On some forms and functions of speech rate changes in everyday conversation. In Auer, P. and Di Luzio, A., editors, *The contextualization of language*, pages 297–336. John Benjamins, Amsterdam.
- Wright, C. E. (1979). Duration differences between rare and common words and their implications for the interpretation of word frequency effects. *Memory & Cognition*, 7(6):411–419.