Language redundancy and acoustic salience: an account in LFG

Tina Bögel
University of Konstanz

Mila Freiseis
University of Konstanz

Romi Hill
Boston University

Daniel Wambach
University of Konstanz

Tianyi Zhao
University of Konstanz

Proceedings of the LFG'24 Conference

Miriam Butt, Jamie Y. Findlay and Ida Toivonen (Editors)

2024

PubliKon

lfg-proceedings.org

Abstract

This paper introduces a new way to integrate gradient language redundancy effects and their acoustic correlates in LFG. Based on data from a production experiment that showed how semantic priming and lexical frequency affect target word duration, the approach models the inverse relationship between language redundancy and acoustic redundancy, with p-structure as a pivotal point between grammar and signal. To this end, all redundancy measures are re-scaled to a common, meaningful scale, while their gradient nature is retained as part of the system. The result not only allows for gradient data to be integrated into the architecture of LFG, but also for the prediction of concrete acoustic measures, thus taking a large step towards the modelling of the phonology-phonetics interface and the generation of spoken language.

1 Introduction

Work in LFG and other frameworks has seen an increased necessity to capture measures of language redundancy in order to account for the preferences speakers have for a particular structure or form (e.g., Bögel 2021; Bresnan 2023).[†] To reflect the probabilities of specific syntactic structures in LFG, for instance, violable ‘soft’ OT(-inspired) constraints can be used to rank several possible structures according to, e.g., syntactic frequency (a.o., Frank et al. 1998) or prosodic input (Butt et al. 2017; Bögel 2020). Similarly, computational approaches have used corpora-based frequency measures to pre-process multiple parses (Cahill et al. 2007). Modelling spoken language, however, remains work in progress due to its often gradient nature, its variability and the lack of reliable and consistent acoustic cues.

One way to approach this problem is to capture speaker preference and the predictability of specific linguistic items by means of language redundancy measures elicited from different modules of grammar, and to relate these measures to specific acoustic realisations (cf. Turk 2010). This relationship between language redundancy and phonetic/phonological parameters has only recently received some attention in LFG. Bresnan (2023), for example, discussed form-reduced verb-pronoun sequences in English (*get them* vs. *get’em*) and proposed Lexical Sharing (Wescoat 2002) to account for the verb-clitic combinations, and Bögel (2021) discussed different pronominal forms in Swabian with respect to lexical i-structure constraints (see Section 4 for a detailed discussion).

In this paper, we propose a new model to account for reduction phenomena based on language redundancy which includes, but also goes beyond the categorical reduction phenomena discussed in Bögel (2021) and Bresnan (2023). Based on the data from a production experiment on semantic priming and lexical frequency, and their effects on durational measures (Freiseis et al. 2024), we show how gradient measures can be modelled in LFG.

[†] This work was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) as part of the DFG-AHRC cooperation with project BO 3388/4-1 *‘Prosodic structure at the interface between language and speech’* in cooperation with Alice Turk and Catherine Lai, University of Edinburgh. We would like to thank the PhonLab, part of the core facilities LingLab at the University of Konstanz, for the use of their recording facilities and the audience and reviewers of LFG’24 for their helpful comments.

The paper is structured as follows: Section 2 outlines the relevant theoretical framework, namely the Smooth Signal Redundancy Hypothesis and the Prosodic Interface Hypothesis. Section 3 presents the production experiment on semantic priming which provided the data for this paper. Section 4 discusses previous LFG-related approaches to language reduction, before focusing on the problem of gradience, the inverse relationship between language redundancy and acoustic salience, and how this relationship can be modelled at p-structure by means of semantic priming and lexical frequency.

2 Language redundancy and acoustic salience: an inverse relationship

This paper follows the *Smooth Signal Redundancy Hypothesis* (henceforth SSRH; Aylett 2000; Aylett & Turk 2004, 2006; Turk 2010), which assumes signal redundancy, i.e., the recognition likelihood of different linguistic items (Turk 2010: 228), to be evenly (smoothly) spread throughout the utterance to ensure robust and efficient communication between interlocutors (cf. Lindblom 1990; Shannon 1948). Signal redundancy is determined by the inverse relationship between a) language redundancy, i.e., the recognition likelihood based on, e.g., lexical, syntactic, semantic or pragmatic factors, and b) acoustic redundancy, i.e., the recognition likelihood based on acoustic salience factors, e.g., duration or fundamental frequency (f_0). Linguistic items with high language redundancy values like frequent words, preferred syntactic structures, or contextually predictable items, would thus result in low acoustic redundancy measures, with shorter durations and smaller pitch/f_0 excursions. Conversely, if language redundancy is low (e.g., as it is with infrequent or unpredictable content words), acoustic salience is predicted to be high, with longer durations and larger pitch excursions (a.o., Bell et al. 2009, 2003; Aylett & Turk 2004, 2006; Bush 2001; Jurafsky et al. 2001; Pluymaekers et al. 2005a,b; Watson et al. 2008; Lam & Watson 2010).

It has further been proposed that this inverse relationship between language redundancy and acoustic salience is mediated via prosodic structure (see Figure 1).

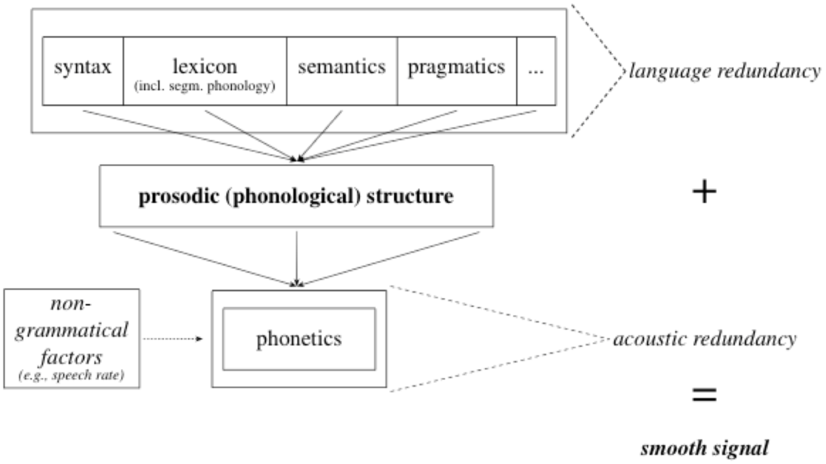

Figure 1: A representation of the SSRH and the PIH (modified from Turk & Shattuck-Hufnagel 2020: 191)

The *Prosodic Interface Hypothesis* (PIH) proposes that language redundancy effects are mostly realised at prosodically prominent sections of a given domain (e.g., word and phrasal stress; Aylett & Turk 2004, 2006), and at prosodic boundaries (Turk 2010), where the boundary-related intervals include the rhyme of the last syllable preceding the boundary and the first onset consonant after the boundary (Fougeron & Keating 1997; Turk & Shattuck-Hufnagel 2007; Dimitrova & Turk 2012). These two target areas in prosodic structure should furthermore be affected by different language redundancy factors in a similar way; i.e., language redundancy factors such as syntactic preferences, lexical frequencies, or bigram frequencies should all be realised by similar acoustic indicators that are associated with the stressed syllables and/or boundary-adjacent intervals, because these are the parts of words that are known to be affected by prosodic prominence and boundaries.

Previous research on the relationship between language redundancy and acoustic salience has been concerned with redundancy effects on the duration of words or morphemes (Fowler & Housum 1987; Pluymaekers et al. 2005a; Bell et al. 2009; Lam & Watson 2010; Kahn & Arnold 2012; Ibrahim et al. 2022b), segmental deletion/lenition/strengthening (Malisz et al. 2018; Brandt et al. 2021; Ibrahim et al. 2022a), also with respect to prosodic boundaries (Bögel & Turk 2019; Andreeva et al. 2020). Investigations above the word level have involved syntactic probabilities (Watson et al. 2006; Levy & Jaeger 2007; Tily et al. 2009), discourse mention, semantic relatedness, and focus (e.g., Lieberman 1963; Balota et al. 1989; Watson & Gibson 2004; Watson et al. 2008; Turnbull 2017).

Recent work by the authors and cooperation partners has been testing the assumptions made by the SSRH and the PIH regarding prosodic prominence and boundary structure in English and German as elicited by duration and f_0 values (Bögel & Turk 2019; Zhang et al. 2023; Zhao et al. 2024; Zhao Forthcoming). In this paper, we present data from an experiment on semantic priming and lexical frequency (Freiseis et al. 2024) and demonstrate how the relationship between language redundancy and acoustic salience (in form of durational measures) can be integrated into the modular structure of LFG.

3 Semantic priming experiment

Semantic priming describes the phenomenon that a word is processed more quickly if preceded by a word that is semantically related. For example, if a subject is presented with the prime *job interview*, they are likely to recognise the target word *applicant* faster than when the target word is preceded by a semantically unrelated word (e.g., *driver*). The semantic priming effect has been attested most notably in various lexical decision tasks, where it was demonstrated that participants have shorter reaction times for primed words than non-primed words (Balota et al. 1989; Foss 1982; Hoedemaker & Gordon 2017; Meyer & Schvaneveldt 1971).

Previous research also revealed an interactive effect between semantic priming and lexical frequency (Becker 1979; Yap et al. 2009; Scaltritti et al. 2013). Becker (1979), following Meyer & Schvaneveldt (1976), argues that semantic priming and increased lexical frequency have a similar activation effect on the target word. As both semantically primed target words and target words with a high lexical frequency are more likely to be recognised, the effect of semantic priming is larger for low frequency words than for high frequency words.

Following the SSRH, target words that are semantically primed are expected to be more redundant, and thus less acoustically salient than non-primed words, and vice versa. Since the SSRH assumes that language redundancy and acoustic salience are mediated by prosodic structure, semantic priming effects are predicted at the boundaries and on the stressed syllable of the target word. The following experiment focuses on the interaction of semantic priming and lexical frequency, and its effect/influence on duration at prosodic word boundaries.[1] A more detailed discussion can be found in Freiseis et al. (2024).

3.1 Methods

3.1.1 Materials

The materials included 22 sentence pairs in Standard German. In each pair, identical target words were presented in either a *priming context* (i.e., where the target word was primed) or a *non-priming context* (i.e., where the target word was not primed). Target words and their lexical frequencies were determined using WebCelex' database (Baayen et al. 2001), and lexical frequency measures for the chosen target words were additionally confirmed using hit numbers of the Google search engine. Table 1 illustrates the thresholds that were used for target words of low and high frequency.

frequency	**WebCelex**	**Google**
high	> 110	> 60 million
low	< 60	< 10 million

Table 1: Thresholds for target words with high or low lexical frequency

The two groups each included six target words which were used for the statistical analysis on the interaction between semantic priming and lexical frequency. The remaining 10 words between these thresholds were disregarded for the statistical analysis on the interaction, but were still included for the general analysis on semantic priming.

Following the predictions made by the SSRH and the PIH, the acoustic effects related to semantic priming and lexical frequency are expected to occur at the boundary-related intervals, i.e., at the interval including the rhyme of the previous word and the onset of the target word, and the interval including the rhyme of the target word and the onset of the following word (see Table 2). In order to guarantee segmentation reliability and comparability of the boundary-related intervals, only target words with plosive onsets were selected. The targets' final syllables were either *-en, -er, -in,* or *-or*. All target words consisted of three syllables and stress on the second syllable. To better distinguish potential priming effects on prosodic boundaries from effects on prosodic

---

[1] The materials used in this study were not suitable for testing effects on the stressed syllable. This would have required segment comparability of the stressed syllable in addition to the comparability of the boundary-related intervals, which makes it difficult to find suitable target words that also fit the frequency requirements. See Zhao et al. (2024) and Zhao (Forthcoming) for a study concerned with both the prominent syllable and prosodic boundaries.

prominence (i.e., the stressed syllable), lexical stress was avoided at the edges of target words.

Target words were common nouns referring to groups of people (e.g., *applicants, pilots*). 18 of the 22 target words appeared in their plural form, ensuring that each target had the same amount of syllables. The target words were preceded by the definite articles *die* (the.FEM/PL) or *der* (the.MASC) and followed by the reflexive pronoun *sich* (‘herself/himself/themselves’) and a verb. The contexts were designed with identical sentence structures and an approximately equal number of syllables. The priming/non-priming context occurred in the first part of the sentence, whereas the second part contained the target word.

Each context sentence pair was used twice with an alternating combination of target words and priming patterns. In the first context sentence, the first part of the sentence was the priming context for the first target word, and the non-priming context for the second target word. For the second context sentence, the order was reversed: The first part of the sentence was the priming context for the second target word, and the non-priming context for the first target word. Example (1) illustrates the material with the target words *Bewerber* ‘applicants’ and *Berliner* ‘people.from.Berlin’. In the first sentence, the priming context in the first part of the sentence is *Vorstellungsgespräch* ‘job interview’, which primes *Bewerber*, but not *Berliner*. In the second sentence, the priming context is *Branderburger Tor* ‘Brandenburg Gate’[2] which primes *Berliner*, but not *Bewerber*. This resulted in two versions for each context sentence in (1) and a total of four sentences for each target word pair.

(1) **Sentence 1: *Bewerber* in a priming context, *Berliner* in a non-priming context**

Um	beim	Vorstellungsgespräch	zu	punkten,	mussten	die	Bewerber/Berliner	sich	behaupten
in.order	at.the	job interview	to	score.points	have.to	the	applicants/Berliner	themselves	assert

‘In order to score points at the job interview, the applicants/people.from.Berlin had to assert themselves.’

Sentence 2: *Berliner* in a priming context, *Bewerber* in a non-priming context

Um	das	Brandenburger Tor	zu	sehen,	mussten	die	Bewerber/Berliner	sich	gedulden
in.order	the	Brandenburg Gate	to	see	have.to	the	applicants/Berliner	themselves	be.patient

‘In order to see the Brandenburg Gate, the applicants/people.from.Berlin had to be patient.’

In order to confirm the priming relationship between context and target word, 18 native speakers of German were asked to check the semantic relatedness via a questionnaire. In the questionnaire, the 22 context sentence pairs were presented and participants were asked to select the intended target from three options (including the primed target, the non-primed target, and an unrelated third word). The intended primed target was chosen in 98.48% of the cases, confirming the semantic relatedness between primes and target words in the materials.

For the production experiment, two experimental lists were created from the 22 context pairs. Each target word only occurred once in each list, either in a priming or a non-priming context, while each context sentence occurred twice, once with a primed target and once with a non-primed target word. The priming context was always presented before the non-priming context, in order to prevent a weakening of the priming context for the primed target and to avoid the creation of a context for the non-primed

---

[2] The Brandenburg Gate is a famous landmark in Berlin.

target. Simultaneously, this design allowed for the weakening of the non-priming context, which - for this experiment - was a desirable effect to ensure that the participant perceived the context as a non-priming one.

3.1.2 Participants

21 German native speakers (mean age = 27, age range 18-30, 15 female and 6 male) participated in the experiment, the majority of whom were students or employees recruited at the University of Konstanz. Participants were randomly assigned to one of the two experimental lists.

3.1.3 Procedure

The recordings took place in a soundproof studio at the University of Konstanz with a condenser microphone (sampling rate 44.1 kHz, 16-Bit, stereo). For the experiment, participants were seated in front of a screen and asked to read out the sentences. After each sentence, the instructor clicked manually to display the next sentence. The procedure took approximately 30 minutes and participants received compensation after the recording.

3.1.4 Analysis

From the resulting 462 sentences, 32 sentences had incorrectly placed lexical stress or showed pronunciation errors and were thus excluded, leaving 430 sentences for the analysis. All sentences were first automatically annotated using MAUS (Kisler et al. 2017). The segmentation was then manually checked and adjusted in accordance with the standard annotation criteria in Turk et al. (2006) using Praat (Boersma & Weenink 2023). Durations of the following six intervals were annotated and extracted: 1) the rhyme of the last syllable of the preceding word (labelled as *R_prev*); 2) the onset of the target word (*O*); 3) the first boundary interval (*B1*) of the target, including the rhyme of the last syllable of the preceding word and the onset of the target; 4) the rhyme of the last syllable of the target (*R*); 5) the second boundary interval (*B2*) of the target word, including the rhyme of the last syllable of the target as well as the onset of the following word, which was the fricative /z/ for all sentences, and 6) the complete target word. Table 2 illustrates a simplified version of the annotation scheme.

(d)ie	*b*	*ewerb*	*er*	*s(ich)*
R_prev	O	–	R	–
B1		–	B2	

Table 2: Annotation scheme for the target word *Bewerber* ‘applicant’ (example (1))

3.2 Results

To assess the effect of semantic priming and lexical frequency on duration measures, linear mixed effects regression models (lmer; Baayen et al. 2008; Kuznetsova et al.

2017; R Core Team 2022) were used, with semantic priming and lexical frequency as fixed factors and participants and items as crossed random factors.[3]

3.2.1 Semantic priming without lexical frequency

A general effect of semantic priming on duration could be established: Primed targets were significantly shorter than non-primed targets in terms of durations of the whole target words (beginning of onset to end of rhyme) ($p < 0.001$). Significant priming effects were also found at the following intervals: 1) Primed targets had a shorter onset (*O*) than their non-primed counterparts ($p < 0.05$), 2) the rhyme of the last target syllable (*R*) was shorter in the primed condition ($p < 0.05$), and 3) the second boundary interval of the targets (*B2*) was shorter in the primed condition as well ($p < 0.05$). In constrast, the first boundary interval *B1* as a whole and the rhyme of the last syllable of the previous word (*R_prev*) did not exhibit any significant priming effect.

3.2.2 Semantic priming effects for frequent and infrequent items

In a second step, the interaction between lexical frequency and semantic priming was calculated. To this end, respective subsets for both frequency and priming conditions were made to take a closer look at the interaction (cf. Section 3.1.1/Table 1). We first started with the semantic priming effects and analysed how the two priming conditions affected frequent and infrequent words, respectively.

For the frequent items, no significant priming effects were found for word-initial intervals. In contrast, for *B2*, primed targets were significantly shorter than non-primed targets ($p < 0.001$). For *R*, the effects approached statistical significance ($p = 0.057$). In contrast, infrequent items were significantly shorter at the beginning of the word (*R_prev, O, B1*) when primed ($p < 0.001$), but not at the end of the word (*R, B2*).

3.2.3 Lexical frequency effects for the primed and the non-primed condition

In order to gain another perspective on the interaction between lexical frequency and semantic priming, a second set of *lmer* models focused on the effects of lexical frequency on primed and non-primed target words.

The onset interval of primed targets yielded a significantly shorter duration for infrequent items than for frequent ones ($p < 0.05$). However, regardless of the priming condition, longer duration was associated with lower lexical frequency towards the end of the targets. Contrary to the results of the target-initial onset, infrequent targets had longer *B2* in both priming ($p < 0.01$) and non-priming condition ($p < 0.05$). The same results were further attested in the sub-interval *R* (primed: $p < 0.01$; non-primed: $p < 0.01$). Figure 2 illustrates the reversed effects of lexical frequency on the onset and the rhyme for the subset of primed target words exclusively.

---

[3] The following section is a condensed version of Freiseis et al. (2024). The reader is referred to this paper for more details.

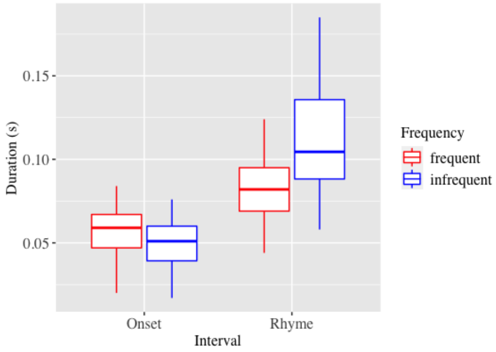

Figure 2: Onset and rhyme duration for the primed data indicating opposite effects of lexical frequency in the word-initial onset and word-final rhyme

3.3 Discussion

The results establish an effect of semantic priming on word duration. The results are also consistent with previous research on the interaction of lexical frequency and semantic priming: The effects of semantic priming seem to be stronger with infrequent items (see also Becker 1979; Yap et al. 2009). However, the results also indicated opposite effects of lexical frequency on the different areas of primed targets (Figure 2), which is not consistent with previous research.

We leave this unexpected observation on the interaction of lexical frequency and primed data for further research (see Freiseis et al. (2024) for some hypotheses). In the remainder of the paper, we will instead focus on the overall priming effects, which were in accordance with the proposals made by the SSRH and the PIH: Semantic priming resulted in shorter durations at the boundary-related intervals of the target words, while non-primed target words featured longer durations. The following sections show, how this gradient relationship between language redundancy factors and acoustic salience can be modelled in LFG.

4 Language redundancy, acoustic salience, and LFG

The results from Section 3 show that language redundancy effects on acoustic salience follow a predictable pattern: As described by the SSRH, there is an inverse relationship between language redundancy and acoustic salience. For our purposes, we can predict that if a lexical item is primed and/or frequent, then duration measures at the boundary-related intervals will be shorter. Vice versa, if a lexical item is non-primed and/or infrequent then its durational measures will be longer. This ensures robust communication between interlocutors, enhancing understanding where needed, and preserving energy where possible.

An open question for LFG (and other frameworks) is how to integrate the regular gradient measures of language redundancy in order to account for speaker preferences and output form. Some previous LFG-related work used violable ‘soft’ OT(-inspired) constraints (Frank et al. 1998) which rank several possibilities to indicate syntactic preferences, also with reference to prosodic indicators (a.o., Butt et al. 2017; Bögel 2020).

However, while LFG-OT constraints are useful to capture the choice between valid syntactic structures, they nevertheless express categorical choices and are thus not suitable to convey real gradience as it is given in the present data and in the relationship between language redundancy and acoustic salience

Recent work by Bresnan (2023) analysed different forms of object pronouns following verbs, i.e., the difference between *get them* and *get'em*. Based on Pierrehumbert's hybrid exemplar-based model of the mental lexicon (Pierrehumbert 2001), Bresnan follows the concept of 'memory traces' of language use. Memory traces include the phonetic forms and their probability distributions learned from the user's experience. Bresnan groups these different forms into two categories, the full form *them* and the clitic *'em*. As some of the verb-clitic combinations exhibit special meanings (e.g., *go get'em* vs. *get/grab them*) and the clitic usually forms a prosodic word with its host, Bresnan proposes Lexical Sharing (Wescoat 2002) as a solution to explain the close relationship between verb and clitic in terms of language redundancy. Under this approach, reduction phenomena and/or phonological cliticisation are explained by assuming that the clitic and its host form one lexical entry projecting to two syntactic nodes, as illustrated in Figure 3.

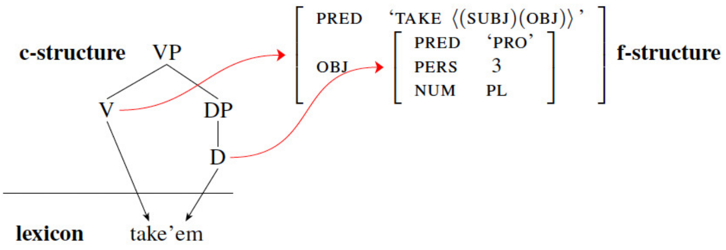

Figure 3: Lexical Sharing as a solution for phonological reduction phenomena (Bresnan 2023: p.69, Figure 3)

The non-compositional meaning of some of the verb-clitic instances indeed points towards lexicalisation. However, this does not imply that all instances of the clitic should be subject to Lexical Sharing, in a similar way that, for example, *bucket* in the idiomatic expression *kick the bucket* does not automatically share a lexical entry with all other items it can occur with. With regard to the argument of prosodic wordhood, it has been established that prosodic words can be smaller or larger than 'lexical' words, and that this mismatch occurs quite frequently with complex morphological structures or function words (a.o., Booij 1984; Nespor & Vogel 1986; Inkelas & Zec 1990; Wheeldon 2000). Consequently, two words forming one prosodic word do not automatically imply Lexical Sharing. In fact, a consequence of Bresnan's approach would be that almost every function word shares its lexical entry with the previous or the following word; a possibility that would lead to infinite additional lexical entries which is not sustainable in terms of efficiency.

An alternative approach to the distinction between pronominal full forms and clitics was discussed in Bögel (2021), where the pronoun retains a separate lexical entry and the prosodic reduction of the clitic form was linked to a) prosodic phrasing constraints, and b) information structural constraints in form of focus and givenness. Under this approach, the two forms reflect (redundancy) information from different modules of the grammar; their final form is determined as part of p-structure. Figure 4 shows the two possible Swabian 1st nominative singular pronoun forms: full form [i] and reduced form [ə].[4] While [i] is a full prosodic word, [ə] can only occur as an enclitic (=σ) similar to the object pronouns discussed in Bresnan (2023). Both forms are associated to the same s-form; their p-forms are distinguished based on whether the pronoun is marked for focus in i-structure, and on prosodic constraints internal to p-structure (e.g., the enclitic cannot occur in the first position of an intonational phrase).

s-form		**p-form**		
PRON	($\uparrow$ PRED) = 'ich'	[i]	SEGMENTS	/i/
	($\uparrow$ PERS) = 1		METR. FRAME	$('\sigma)_{\omega}$
	($\uparrow$ CASE) = nom	- - -	- - -	- - -
	($\uparrow$ PRON) = pers	[ə]	SEGMENTS	/ə/
	($\uparrow$ NUM) = sg		METR. FRAME	$=\sigma$
	...		$\neg(\uparrow_i$ FOCUS)	

Figure 4: Lexical entries for the Swabian 1NomSg pronoun [i] and its reduced form [ə], modified from (Bögel 2021: 19)

This approach incorporates the well documented relationship between form reduction and language redundancy constraints expressed by the SSRH and the PIH discussed in Section 2: Prosodic structure regulates the inverse relationship between language redundancy and acoustic salience. If language redundancy is low (i.e., the pronoun is in focus), acoustic salience is high (i.e., the full form is applied) and if language redundancy is high (i.e., the pronoun in its unmarked context) acoustic salience is low and the reduced form is applied.

However, neither of these approaches, Bögel (2021) and Bresnan (2023), can account for gradient data. Pronouns are not considered to be gradient, but they usually are assumed to encompass more than the two forms discussed above. For the pronoun *them*, for example, this would include *them, 'em, 'm,* and ϕ, where the phonological realisation has been tied to language redundancy, namely a 'givenness' scale (a.o., Baumann 2006). This categorical scale of four forms already poses a problem for Lexical Sharing, but could be integrated into the approach illustrated in Figure 4. Problematic for both are the numerous cases where surface forms are reduced on gradient scales and thus do not allow for a natural categorical classification. This paper tackles this problem and introduces a new proposal to integrate gradient data into LFG, following the assumptions made by the SSRH and the PIH on the inverse relationship between language redundancy and acoustic salience, with prosodic structure as the pivotal point between them.

4.1 Gradient data is expressed by gradient scales

The experiment discussed in Section 3 is concerned with two different language redundancy factors: semantic priming and lexical frequency. For the statistical analysis, the

---

[4] See Section 4.3. for a brief overview on the multidimensional lexicon as proposed in Bögel (2015).

two factors were treated as categorical: primed and non-primed, and high and low frequency. However, this classification ignores the fact that both groups are also gradient within the respective categories. Furthermore, the SSRH does not predict one duration for all elements in one category, but a successive gradation in relation to the corresponding language redundancy measures.

In order to obtain representative scales for the lexical frequencies and the semantic priming measures of the experimental target words, we extended the methods that were used to create the materials (see Section 3.1.1). For the raw lexical frequencies of our target words, we used the German Wikipedia dataset of *Projekt Deutscher Wortschatz* (PDW; Goldhahn et al. 2012).[5] In order to sum lexical frequencies across inflected forms, we tokenized the lemmatized dataset with SpaCy (Montani et al. 2023). For stop word removal, we use a modified version of a stop word list for German (Savand et al. 2024).[6]

Within the 5 million Wikipedia sentences, there was a great variability between the raw word frequencies. For example, *benehmen* 'to behave' occurs only 55 times, while *sehen* 'to see' has a frequency count of 52205. Such differences will only increase with larger corpora, with many infrequent words having low numbers and only a few frequent words with high numbers (following Zipf's law; Zipf 1936, 1949); a problem that requires a solution before the data can be used in our model.

To measure semantic priming we used co-occurrence counts and positive pointwise mutual information (PPMI) (Fano 1961; Church & Hanks 1990), which is high for words that are closely related to each other and low for unrelated tokens, taking the general frequency of a word and its likelihood to co-occur with other words into account as well. The formula is provided in the equation in (2), where the PPMI measures the association between two words x and y. The probability of the words co-occurring is divided by the probability of each word occurring independently. The denominator thus takes into account the lexical frequencies of the individual words. Closely related words are those that co-occur more often than predicted under independence.

(2) $$PPMI(x,y) = max(log2\frac{P(x,y)}{P(x)P(y)}, 0).$$

Word frequencies and co-occurrence counts are required for the denominator and numerator respectively. Both of these calculations were extracted from the German Wikipedia dataset of the PDW with the same pre-processing techniques described above for lexical frequencies. The co-occurrences were calculated at the sentence level. To do so, a data frame was created which stored all the words in the corpus, their respective co-occurring tokens and the number of co-occurrences for each of these pairs. Based on the frequency counts and this data frame, we calculated the PPMI score between any two words. Similar to the Zipf distribution for lexical frequencies, the PPMI score distribution is also skewed. Most words do not co-occur with each other, resulting in 0.0

---

[5] Lexical frequencies do not necessarily represent a specific participant's familiarity with different words, but they provide a good approximation for a majority of the population.

[6] This particular list was chosen and extended because the ones offered by common libraries like SpaCy or NLTK (Bird et al. 2009) are too extensive for our cause, and would exclude words that we deem important for the calculations.

PPMI scores for the majority of word pairs.

As mentioned above, we only focus on onset closure durations (O-values) of the word-initial plosives from the experiment described in Section 3 for demonstration purposes, but the method can easily be extended to all other target areas. For the onset duration scale we extracted all onset measures from the experimental production data, excluding any outliers beyond the interquartile range.

For the research question discussed in this paper we thus have three gradient scales that express different measures of language redundancy and acoustic salience:

- lexical frequencies: 0 to millions of occurrences
- semantic priming: 0 to millions of co-occurrences
- onset closure duration: 0.02 to 0.102 seconds

With other language redundancy factors, different target areas, or further acoustic cues (e.g., f_0 in Hertz instead of duration in seconds), there are potentially many more scales expressing the gradient nature of the data, with many of them being non-linear. There are thus two problems that need to be solved with regard to gradient data: 1) How can we relate the different scales of - in our case - semantic priming in form of PPMI scores, lexical frequencies, and onset closure durations? And 2) how can we meaningfully integrate these measures into a formal model of grammar?

4.2 Normalising and modelling the inverse relation

In order to curtail the non-linear scales and to map these different measures onto a single scale, we applied logarithmic normalisation to the raw occurrence counts, duration values, and PPMI scores. The log-transformation represents all measures on a scale between 0 and 1, while still capturing the gradient nature of the data, even if the original data is heavily skewed (lexical frequencies and PPMI scores). In order to make all scales comparable, we apply the log-transformation also to the duration values, even though they are normally distributed. Example (3) shows the corresponding formula.

(3) $$\text{log normalisation} = \frac{\ln(\text{value}) - \ln(\text{minimum})}{\ln(\text{maximum}) - \ln(\text{minimum})}$$

For the language redundancy factors of lexical frequency and semantic priming, this subsequently means that values close to 1 represent a high language redundancy value. Values close to 0 on the other hand indicate a low redundancy. For duration values, scores close to 0 represent short durations, i.e., low acoustic salience, and numbers close to 1 represent longer durations and thus high acoustic salience. The interpretation of the scales is thus reversed: High scores on the language redundancy scale represent frequent items, while high scores on the acoustic redundancy scale represent infrequent items.

In order to model the inverse relationship between language redundancy and acoustic salience, and to predict the acoustic redundancy score during language production, we subtract the log-normalised language redundancy value from 1 (the maximum log-value an item can have) and receive the log-normalised acoustic redundancy value. This log-value can then be mapped back against the concrete acoustic values from the experiment, in our case onset closure duration.

(4) 1– language redundancy (log) = acoustic redundancy (log)
→ onset closure duration (prediction in s)

To give a concrete example: The word *Berliner* has a raw lexical frequency value of 9610 occurrences which is normalised to the corresponding log-value of 0.827. Subtracted from 1, the returned inverse acoustic redundancy value will be 0.173.

(5) 1 – 0.827 = 0.173

In a second step, the acoustic redundancy log-value is mapped against concrete onset measures. As mentioned above, the onset closure durations collected as part of the experiment were log-normalised as well. The log-value from example (5) corresponds to the concrete value of 0.027 seconds in the onset closure duration data.

(6) 0.173 log-value ≈ 0.027 onset closure duration in seconds

The onset closure duration is rather short in this example, reflecting the high lexical frequency of the word *Berliner*. Of course, it is unlikely that the predicted exact value of 0.027 seconds is produced by a speaker uttering the word *Berliner*. There are many factors influencing spoken language, with lexical frequency just being one of them. However, we can predict that based on lexical frequency, the duration of onset closure duration will be *approximately* 0.027 seconds.

To summarize, our approach allows us to transform the different raw value scales of lexical frequency, semantic priming, and onset closure duration to a log-transformed scale between 0 and 1. For language redundancy values, items with the value of 1 are highly frequent/predictable and items with the value of 0 are infrequent/non-predictable (with values in between trending towards one end of the scale). For duration measures, the scales are reversed, in that a value of 1 represents a long onset closure duration, and a value of 0 a short duration. The normalised scales then allow us to model the inverse relationship between language redundancy and acoustic redundancy as illustrated in Table 3.

word	raw freq.	language red.		acoustic red.	onset clos.dur
Berliner	9610	0.827	1 – 0.827	0.173	0.027

Table 3: Modelling the inverse relationship between language redundancy and acoustic redundancy by means of normalised log-scales for the word *Berliner*

This approach allows us to effectively calculate realistic acoustic values based on language redundancy values. The next sections discuss how these values can be modelled in LFG with p-structure providing the pivotal point between language redundancy and acoustic redundancy as predicted by the PIH.

4.3 Recap: The prosody-syntax interface (Bögel 2015)

The new proposal introduced in this paper is based on the syntax-prosody interface model developed in Bögel (2015), which assumes two levels of information exchange between c-structure and p-structure: 1) The transfer of vocabulary which exchanges phonological and morphosyntactic information of lexical elements via the multidimensional lexicon, and 2) the transfer of structure (♮) which exchanges information on syntactic and prosodic constituency, and on intonation. The model distinguishes between comprehension (i.e., parsing; from form to meaning) and production (i.e., generation; from meaning to form), slightly altering communication at the interface in each case.[7] During comprehension, information from the speech signal feeds into p-structure; during production, information from p-structure forms the basis for the speech signal/the utterance. Figure 5 illustrates.

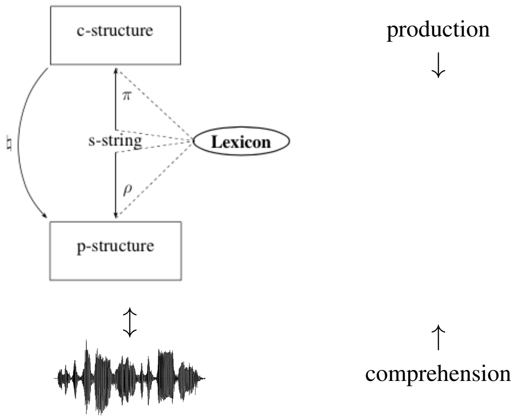

Figure 5: The prosody-syntax interface as proposed in Bögel (2015)

The multidimensional lexicon associates morphosyntactic (s-form) and phonological information (p-form) on lexical elements and projects them to their respective structures (c-structure or p-structure). Figure 6 shows the (simplified) lexical entries for the noun *Bewerber* 'applicant' and the determiner *die* 'the'.

concept	**s-form**			**p-form**	
APPLICANT	N	(↑ PRED)	= 'Bewerber'	SEGMENTS	/b ə v ɛ ʁ b ɐ/
		(↑ NUM)	= {sg | pl}	METRICAL FRM	$(\sigma'\sigma\sigma)_\omega$
DETERMINER	D	(↑ PRED)	= 'die'	SEGMENTS	/d i/
		(↑ NUM)	= {sg | pl}	METRICAL FRM	σ

Figure 6: (Simplified) lexical entries for *die* and *Bewerber*

Next to the (semantic) concept and the morpho-syntactic s-form, each lexical entry also contains information on its phonological representation in p-form. This includes the segments and the metrical frame, which contains the number of syllables, information on lexical stress, and whether the entry itself is a prosodic word (*Bewerber*) or not (*die*).

---

[7] See Bögel (2020) for a detailed example, and Bögel (2023) for a discussion.

Each lexical dimension can only be accessed by the related module, i.e., p-structure can only access information from the p-form, and c-structure can only access information from the s-form. At the same time, the lexicon also has a translation function: Once a dimension is triggered, the related dimensions can be accessed as well. During comprehension, the input from the speech signal is matched against the lexicon's p-form until a match is made and the s-form becomes available for processing. During production, information on a particular s-form is passed from the grammar to the lexicon, and the associated p-form becomes available to p-structure.

The model proposed in this paper is discussed from the perspective of production, i.e., how semantic priming and lexical frequency effects are realised as part of the acoustic signal. P-structure is represented by the p-diagram, a syllablewise representation of the speech signal over time. During production, the p-diagram consists of two levels, the *lexical* level with information from the lexicon's p-form, and the *interpretation* level, which, e.g., contains information on c-structure constituents translated into prosodic constituents. The p-diagram in Figure 7 contains the string ... *mussten die Bewerber* ... from example (1).

PHRASING	$(_{\iota}(\sigma$	$\sigma)_{\omega}$	σ	$_{\omega}(\sigma$	σ	$\sigma)_{\omega}$	...	*interpretation*
TOBI	–	–	–	L	+H*	–	...	↓
LEX.STRESS	x	–	–	–	x	–	...	*lexical*
SEGMENTS	/mʊs/	/tən/	/di/	/bə/	/vɛʁ/	/bɐ/	...	↓
VECTORINDEX	$\mathbf{S_1}$	$\mathbf{S_2}$	$\mathbf{S_3}$	$\mathbf{S_4}$	$\mathbf{S_5}$	$\mathbf{S_6}$	...	

Figure 7: A (simplified) p-diagram during production for the string ... *mussten die Bewerber* ... (example (1))

For each word, the p-diagram stores the syllables, the segments, lexical stress, and prosodic frame as defined in the lexicon. In addition, further information on high and low tones and on larger prosodic units (here the start of an intonational phrase ($(_{\iota}$)) are added based on the lexical information and information from the grammar.[8] This initial p-diagram includes the basic phonological information, which is then further adjusted according to language specific prosodic rules and principles. However, so far – and in contrast to the model in comprehension – the production model does not yet include an interface to phonetics (i.e., the actual speech signal), and it does not provide a means to express language redundancy and its inverse relationship with acoustic salience.

4.4 Modelling language redundancy in LFG: a new proposal

Section 4.2 showed how different factors of language redundancy and acoustic redundancy can be normalised on one scale, and how their inverse relationship can be expressed. This section demonstrates how redundancy measures can be integrated into LFG and how p-structure can be used as a pivot between grammar and signal, using the proposal of Bögel (2015) as discussed in Section 4.3.

---

[8] Further information is left out in order to simplify the discussion.

In order to integrate lexical frequencies, the multidimensional lexicon is extended to include a *redundancy* section as illustrated in Figure 8 for the words *Berliner, Bewerber* and *Vorstellungsgespräch* from example (1).[9]

s-form		p-form		redundancy
N	(↑ PRED) = 'people.from.Berlin' (↑ NUM) = pl	P-FORM METR. FRAME	[bɛʁli:nɐ] $(\sigma'\sigma\sigma)_\omega$	lex.freq: 0.827

s-form		p-form		redundancy
N	(↑ PRED) = 'applicants' (↑ NUM) = pl	P-FORM METR. FRAME	[bəvɛʁbɐ] $(\sigma'\sigma\sigma)_\omega$	lex.freq: 0.579

s-form		p-form		redundancy
N	(↑ PRED) = 'job interview' (↑ NUM) = sg	P-FORM METR. FRAME	[fo:ɐʃtɛlʊŋsgəʃpʁɛ:ç] $('\sigma\sigma\sigma'\sigma)_\omega$	lex.freq: 0.18

Figure 8: Lexical entries for the words *Berliner, Bewerber* and *Vorstellungsgespräch*, including log-transformed lexical frequency values

These values become available to p-structure during the transfer of vocabulary. To this end, the p-structure is extended to include a *redundancy level* which stores the lexical frequency value for all the syllables associated with that lexical entry. This input to p-structure is further processed to calculate the corresponding measure of acoustic salience. Figure 9 demonstrates this process for the word *Bewerber* 'applicant'.

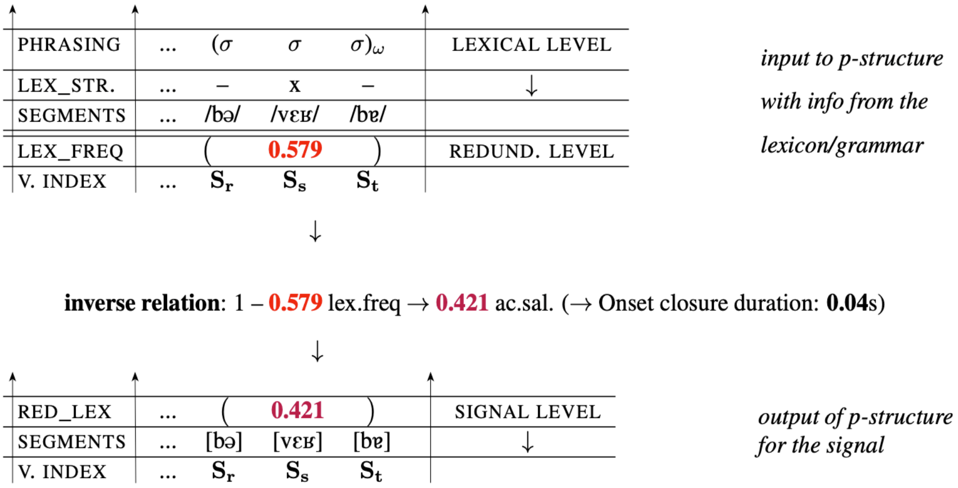

PHRASING	...	$(\sigma$	σ	$\sigma)_\omega$	LEXICAL LEVEL
LEX_STR.	...	–	x	–	↓
SEGMENTS	...	/bə/	/vɛʁ/	/bɐ/	
LEX_FREQ		(	**0.579**	)	REDUND. LEVEL
V. INDEX	...	$\mathbf{S_r}$	$\mathbf{S_s}$	$\mathbf{S_t}$	

input to p-structure
with info from the
lexicon/grammar

↓

inverse relation: 1 – **0.579** lex.freq → **0.421** ac.sal. (→ Onset closure duration: **0.04**s)

↓

RED_LEX	...	(	**0.421**	)	SIGNAL LEVEL
SEGMENTS	...	[bə]	[vɛʁ]	[bɐ]	↓
V. INDEX	...	$\mathbf{S_r}$	$\mathbf{S_s}$	$\mathbf{S_t}$	

output of p-structure
for the signal

Figure 9: Prosodic structure as the pivotal point between language redundancy and acoustic redundancy for *Bewerber* 'applicant'

Once the log-value of the acoustic salience is established, it can be mapped to a concrete acoustic value as part of the p-structure-phonetics interface. This can be the concrete onset closure duration (which in this case would be 0.04 s), but it could also be mapped against any other acoustic value associated with the target word (overall and further partial duration values, f_0 values, etc.). The acoustic salience measure thus forms an

---

[9] The redundancy section in this example only contains the lexical frequency measures. Further measures could, e.g., be syllable bigrams which have also been shown to have an effect on acoustic salience.

abstract representation for a number of phonetic realisations while at the same time allowing for their gradient nature to be captured as well.

Semantic priming is different from lexical frequency as it does not express the activation of an isolated word, but the relationship between two different words. This network between words is also part of the mental lexicon, but not part of a single lexical entry's redundancy level. Figure 10 shows the semantic priming measures (i.e., the log-normalised PPMI values from the corpus) between single lexical entries in blue. The numbers indicate a strong priming relationship between 'applicant' and 'job interview' (0.963) and no priming relationship between 'job interview' and 'Berliner' (0). In p-structure, semantic priming measures are added to the new redundancy level.[10] As it is the case with lexical frequency, semantic priming measures can be inversely correlated with the log-value for acoustic redundancy (0.037 in Figure 10).

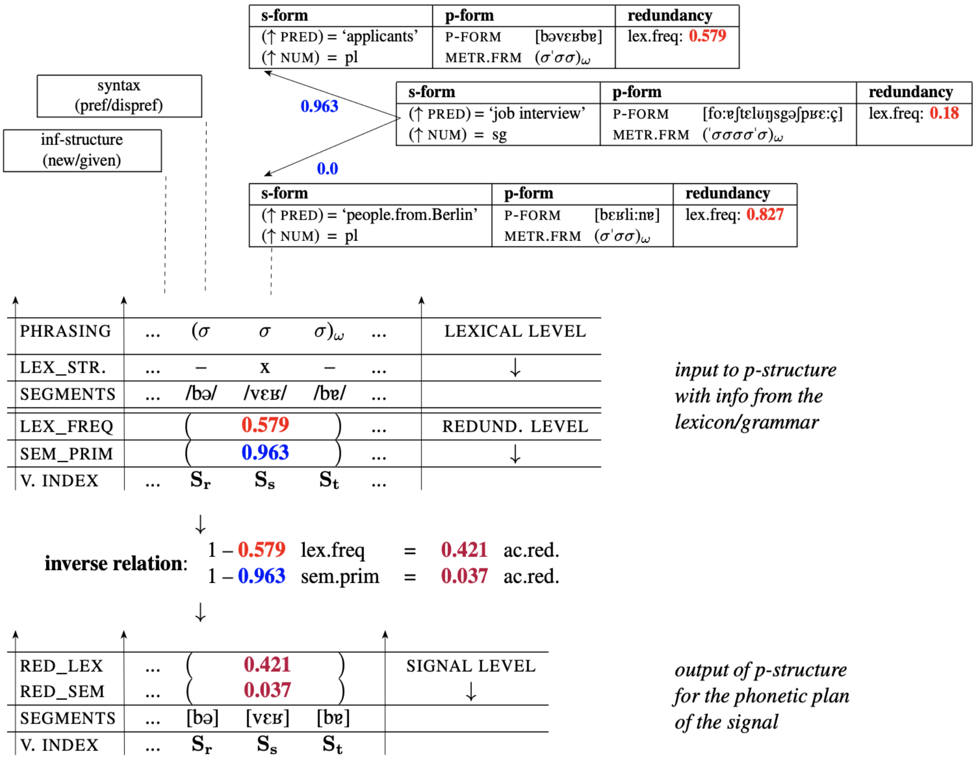

Figure 10: Calculated measures of semantic priming, lexical frequency, and acoustic redundancy for the word *Bewerber* at the prosody-syntax interface Bögel (2015)

The remaining question is how these two acoustic redundancy measures determined by semantic priming (0.421) and lexical frequency (0.037) can be united. Is the mean between the two values an appropriate representation? Or does one value 'override' the other value? Initial results as discussed in Section 3.2 show that the relationship is rather complex, so we will leave this question for further research.

---

[10] Further values to add to the redundancy level could be, for example, syntactic preferences or information from i-structure on focused or given structures.

5 Conclusion

This paper introduced a new approach to integrate gradient language redundancy effects into a formal model of grammar. Following the Smooth Signal Redundancy Hypothesis and the Prosodic Interface Hypothesis, the approach assumes that language redundancy and acoustic redundancy form an inverse relationship mediated by prosodic structure. If language redundancy is high, acoustic redundancy is low, and vice versa, with the effects mainly being realised at the prosodic boundaries and on the stressed syllable.

The concept was demonstrated by means of an experiment on semantic priming and lexical frequency and their inverse relationship with durational values at the prosodic word boundaries as an acoustic redundancy measure. To model this relationship, these three types of gradient redundancy measures were log-normalised to a common scale between 0 and 1. The normalisation allowed for an abstract representation of the different redundancy values, while at the same time preserving the gradient nature of the data. The common scale furthermore made it possible to model the inverse relationship between language redundancy and acoustic redundancy measures. These gradient representations were then integrated into LFG, with p-structure as the pivotal point between language redundancy and acoustic salience.

The formal integration of gradient redundancy measures in a way presupposes the assumption that these measures are part of the (rule-based) grammar, which opens up an interesting discussion on the exact definition of a speaker's language competence vs. performance. While this discussion is beyond the scope of this paper, the paper showed that the relationship between language redundancy and acoustic salience can be captured by regular inverse correlation patterns. The integration of gradient data also avoids the commonly found classification of such values into different categories based on random thresholds (e.g., 'frequent' and 'infrequent', or 'given' and 'new'). Instead, the proposal is able to preserve the original nature of the data on the one hand, and the fundamental rule-based structure of LFG on the other hand, by adding redundancy values to already existing structures and by providing a predictive algorithm to calculate specific redundancy values.

In addition, the paper also takes an important step towards the interface between phonology and phonetics by means of abstract acoustic redundancy values. These values can be further transformed into concrete acoustic measures, thus providing a means to support the generation of spoken language based on deep linguistic information as it is traditionally found in LFG.

Appendix

The following section lists the stimuli used in the experiment described in Section 3. The sentence pairs are grouped with identical target words in a priming (first sentence) and a non-priming (second sentence) context. All target words are highlighted in blue.

1. Um die Nachbarn zu erfreuen, mussten die Bewohner sich benehmen.
   Um die Botschaft zu verkünden, mussten die Bewohner sich bemühen.

2. Um die Botschaft zu verkünden, mussten die Propheten sich bemühen.
   Um die Nachbarn zu erfreuen, mussten die Propheten sich benehmen.

3. Um das Brandenburger Tor zu sehen, mussten die Berliner sich gedulden.
   Um beim Vorstellungsgespräch zu punkten, mussten die Berliner sich behaupten.
4. Um beim Vorstellungsgespräch zu punkten, mussten die Bewerber sich behaupten.
   Um das Brandenburger Tor zu sehen, mussten die Bewerber sich gedulden.
5. Um das Eigentum zu behalten, mussten die Besitzer sich beweisen.
   Um den Eiffelturm zu besuchen, mussten die Besitzer sich gedulden.
6. Um den Eiffelturm zu besuchen, mussten die Pariser sich gedulden.
   Um das Eigentum zu behalten, mussten die Pariser sich beweisen.
7. Um das Geld zu klauen, mussten die Banditen sich verbünden.
   Um das Haar zu föhnen, mussten die Banditen sich beeilen.
8. Um das Haar zu föhnen, mussten die Blondinen sich beeilen.
   Um das Geld zu klauen, mussten die Blondinen sich verbünden.
9. Um die Römer zu besiegen, mussten die Germanen sich verbünden.
   Um den Bruder zu verärgern, mussten die Germanen sich bemühen.
10. Um den Bruder zu verärgern, mussten die Geschwister sich bemühen.
    Um die Römer zu besiegen, mussten die Geschwister sich verbünden.
11. Um die Ehe zu retten, musste die Gemahlin sich bemühen.
    Um den Gegner zu schlagen, musste die Gemahlin sich beweisen.
12. Um den Gegner zu schlagen, musste der Gewinner sich beweisen.
    Um die Ehe zu retten, musste der Gewinner sich bemühen.
13. Um das Königreich zu walten, musste die Prinzessin sich benehmen.
    Um die Vorlesung zu halten, musste die Prinzessin sich beweisen.
14. Um die Vorlesung zu halten, musste der Professor sich beweisen.
    Um das Königreich zu walten, musste der Professor sich benehmen.
15. Um das Frachtschiff zu kapern, mussten die Piraten sich verbünden.
    Um das Flugzeug zu landen, mussten die Piraten sich beeilen.
16. Um das Flugzeug zu landen, mussten die Piloten sich beeilen.
    Um das Frachtschiff zu kapern, mussten die Piloten sich verbünden.
17. Um die Sehenswürdigkeit zu betrachten, mussten die Touristen sich gedulden.
    Um an der Universität zu lehren, mussten die Touristen sich beweisen.
18. Um an der Universität zu lehren, mussten die Dozenten sich beweisen.
    Um die Sehenswürdigkeit zu betrachten, mussten die Dozenten sich gedulden.
19. Um das Volk zu unterdrücken, mussten die Tyrannen sich verbünden.
    Um den Chef zu beeindrucken, mussten die Tyrannen sich behaupten.
20. Um den Chef zu beeindrucken, mussten die Kollegen sich behaupten.
    Um das Volk zu unterdrücken, mussten die Kollegen sich verbünden.

21. Um den Gefährten zu helfen, mussten die Kumpanen sich vertrauen.
    Um den Eintritt zu bezahlen, mussten die Kumpanen sich gedulden.

22. Um den Eintritt zu bezahlen, mussten die Besucher sich gedulden.
    Um den Gefährten zu helfen, mussten die Besucher sich vertrauen.

References

Andreeva, Bistra, Bernd Möbius & James Whang. 2020. Effects of surprisal and boundary strength on phrase-final lengthening. In *Proceedings of the 10th International Conference on Speech Prosody 2020*, 146–150. https://www.coli.uni-saarland.de/~andreeva/PAPERS/Effects_of_surprisal_and_boundary_strength_on_phra.pdf.

Aylett, Matthew. 2000. *Stochastic Suprasegmentals: Relationships between Redundancy, Prosodic Structure and Care of Articulation in Spontaneous Speech*. Ph.D. thesis, University of Edinburgh. http://hdl.handle.net/1842/22455.

Aylett, Matthew & Alice Turk. 2004. The Smooth Signal Redundancy Hypothesis: a functional explanation for relationships between redundancy, prosodic prominence, and duration in spontaneous speech. *Language and Speech* 47(1). 31–56. https://doi.org/10.1177/00238309040470010201.

Aylett, Matthew & Alice Turk. 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. https://doi.org/https://doi.org/10.1121/1.2188331.

Baayen, R. Harald, Doug J. Davidson & Douglas M. Bates. 2008. Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language* 59(4). 390–412. https://doi.org/https://doi.org/10.1016/j.jml.2007.12.005.

Baayen, R.H., R. Piepenbrock & L. Gulikers. 2001. *WebCelex*. Online resource: http://celex.mpi.nl/. Max Planck Institute for Psycholinguistics.

Balota, David A., Julie E. Boland & Lynne W. Shields. 1989. Priming in pronunciation: Beyond pattern recognition and onset latency. *Journal of Memory and Language* 28(1). 14–36. https://doi.org/https://doi.org/10.1016/0749-596X(89)90026-0.

Baumann, Stefan. 2006. *The Intonation of Givenness* (Linguistische Arbeiten 508). Tübingen: Niemeyer.

Becker, Curtis A. 1979. Semantic context and word frequency effects in visual word recognition. *Journal of Experimental Psychology: Human Perception and Performance* 5(2). 252. https://doi.org/https://doi.org/10.1037/0096-1523.5.2.252.

Bell, Alan, Jason M. Brenier, Michelle Gregory, Cynthia Girand & Dan Jurafsky. 2009. Predictability effects on durations of content and function words in conversational English. *Journal of Memory and Language* 60(1). 92–111. https://web.stanford.edu/~jurafsky/bell09.pdf.

Bell, Alan, Daniel Jurafsky, Eric Fosler-Lussier, Cynthia Girand, Michelle Gregory & Daniel Gildea. 2003. Effects of disfluencies, predictability, and utterance position on word form variation in English conversation. *Journal of the Acoustical Society of America* 113(2). 1001–1024. https://doi.org/10.1121/1.1534836.

Bird, Steven, Edward Loper & Ewan Klein. 2009. *Natural Language Processing with Python*. O'Reilly Media Inc. https://github.com/nltk/nltk.

Boersma, Paul & David Weenink. 2023. Praat: doing phonetics by computer [computer program, Version 6.4.01]. Available at http://www.praat.org/.

Bögel, Tina. 2015. *The Syntax–Prosody Interface in Lexical Functional Grammar*. Ph.D. thesis, University of Konstanz. https://kops.uni-konstanz.de/entities/publication/e4ff5b6d-e028-4d52-9bce-fb7367035e41.

Bögel, Tina. 2020. German case ambiguities at the interface: production and comprehension. In Gerrit Kentner & Joost Kremers (eds.), *Prosody in Syntactic Encoding* (Linguistische Arbeiten 573), 51–84. Berlin: De Gruyter. https://doi.org/10.1515/9783110650532.

Bögel, Tina. 2021. Function words at the interface: A two-tier approach. *Languages* 6(4). 197.

Bögel, Tina. 2023. Prosody and its interfaces. In Mary Dalrymple (ed.), *The Handbook of Lexical Functional Grammar*, Berlin: Language Science Press. https://doi.org/10.5281/zenodo.10185970.

Bögel, Tina & Alice Turk. 2019. Frequency effects and prosodic boundary strength. In S. Calhoun, P. Escudero, M. Tabain & P. Warren (eds.), *Proceedings of the International Congress of Phonetic Sciences (ICPhS)*, 1014–1018. https://assta.org/proceedings/ICPhS2019/papers/ICPhS_1063.pdf.

Booij, Geert E. 1984. Principles and parameters in prosodic phonology. In Brian Butterworth, Bernard Comrie & Östen Dahl (eds.), *Explanations for Language Universals*, 249–280. Berlin, Boston: De Gruyter Mouton. https://doi.org/doi:10.1515/9783110868555.249.

Brandt, Erika, Bernd Möbius & Bistra Andreeva. 2021. Dynamic formant trajectories in German read speech: Impact of predictability and prominence. *Frontiers in Communication* 6. 643528. https://doi.org/https://doi.org/10.3389/fcomm.2021.643528.

Bresnan, Joan. 2023. Why lexical syntax? Evidence from English object pronoun enclitics. In *Proceedings of LFG23*, 66–89. Konstanz: PubliKon. https://web.stanford.edu/~bresnan/bresnan-lfg23proceedings.pdf.

Bush, Nathan. 2001. Frequency effects and word-boundary palatalization in English. In Joan Bybee & Paul Hopper (eds.), *Frequency and the emergence of linguistic structure*, Amsterdam: John Benjamins. https://doi.org/https://doi.org/10.1075/tsl.45.14bus.

Butt, Miriam, Tina Bögel & Farhat Jabeen. 2017. Polar *kya* and the prosody-syntax-pragmatics interface. In *Proceedings of LFG'17*, Stanford, CA: CSLI Publications.

Cahill, Aoife, Tracy Holloway King & John T. Maxwell III. 2007. Pruning the search space of a hand-crafted parsing system with a probabilistic parser. In Timothy Baldwin, Mark Dras, Julia Hockenmaier, Tracy Holloway King & Gertjan van Noord (eds.), *ACL 2007 workshop on deep linguistic processing*, 65–72. Prague, Czech Republic: Association for Computational Linguistics. https://aclanthology.org/W07-1209.

Church, Kenneth Ward & Patrick Hanks. 1990. Word association norms, mutual information, and lexicography. *Computational Linguistics* 16(1). 22–29. https://doi.org/https://aclanthology.org/J90-1003.

Dimitrova, Snezhina & Alice Turk. 2012. Patterns of accentual lengthening in English four-syllable words. *Journal of Phonetics* 40(3). 403–418. https://doi.org/https://doi.org/10.1016/j.wocn.2012.02.008.

Fano, Robert. 1961. *Transmission of Information: A Statistical Theory of Communication*. MIT Press. OCLC: 1313256906.

Foss, Donald. 1982. A discourse on semantic priming. *Cognitive Psychology* 14. 590–607. https://doi.org/10.1016/0010-0285(82)90020-2.

Fougeron, Cécile & Patricia A. Keating. 1997. Articulatory strengthening at edges of prosodic domains. *Journal of the Acoustical Society of America* 101. 3728–3740. https://doi.org/10.1121/1.418332.

Fowler, Carol A. & Jonathan Housum. 1987. Talkers' signaling of "new" and "old" words in speech and listeners' perception and use of the distinction. *Journal of Memory and Language* 26(5). 489–504. https://doi.org/https://doi.org/10.1016/0749-596X(87)90136-7.

Frank, Anette, Tracy Holloway King, Jonas Kuhn & John T. Maxwell III. 1998. Optimality Theory style constraint ranking in large-scale LFG grammars. In *Proceedings of LFG'98*, Stanford, CA: CSLI Publications. https://doi.org/https://web.stanford.edu/group/cslipublications/cslipublications/LFG/LFG3-1998/lfg98franketal.pdf.

Freiseis, Mila, Tianyi Zhao & Tina Bögel. 2024. Semantic priming and prosodic structure: At the interface between language redundancy and acoustic salience. In *Proccedings of Speech Prosody 2024*, 1015–1019. Leiden, the Netherlands. https://doi.org/10.21437/SpeechProsody.2024-205.

Goldhahn, Dirk, Thomas Eckart & Uwe Quasthoff. 2012. Building large monolingual dictionaries at the Leipzig corpora collection: From 100 to 200 languages. In Nicoletta Calzolari, Khalid Choukri, Thierry Declerck, Mehmet Uğur Doğan, Bente Maegaard, Joseph Mariani, Asuncion Moreno, Jan Odijk & Stelios Piperidis (eds.), *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, 759–765. Istanbul, Turkey: European Language Resources Association (ELRA). http://www.lrec-conf.org/proceedings/lrec2012/pdf/327_Paper.pdf".

Hoedemaker, Renske & Peter Gordon. 2017. The onset and time course of semantic priming during rapid recognition of visual words. *Journal of Experimental Psychology: Human Perception and Performance* 43. https://doi.org/10.1037/xhp0000377.

Ibrahim, Omnia, Ivan Yuen, Bistra Andreeva & Bernd Möbius. 2022a. The effect of predictability on German stop voicing is phonologically selective. In *Proceedings of Speech Prosody 2022*, 669–673. https://doi.org/10.21437/SpeechProsody.2022-136.

Ibrahim, Omnia, Ivan Yuen, Marjolein van Os, Bistra Andreeva & Bernd Möbius. 2022b. The combined effects of contextual predictability and noise on the acoustic realisation of German syllables. *The Journal of the Acoustical Society of America* 152(2). 911–920. https://doi.org/https://doi.org/10.1121/10.0013413.

Inkelas, Sharon & Draga Zec (eds.). 1990. *The Phonology-Syntax Connection*. CSLI Publications.

Jurafsky, Dan, Alan Bell, Michelle Gregory & William Raymond. 2001. Probabilistic relations between words: Evidence from reduction in lexical production. In Joan Bybee & Paul Hopper (eds.), *Frequency and the Emergence of Linguistic Structure*, 229–254. Amsterdam: John Benjamins.

Kahn, Jason M. & Jennifer E. Arnold. 2012. A processing-centered look at the contribution of of givenness to durational reduction. *Journal of Memory and Language* 67. 311–325. https://doi.org/https://doi.org/10.1016/j.jml.2012.07.002.

Kisler, Thomas, Uwe Reichel & Florian Schiel. 2017. Multilingual processing of speech via web services. *Computer Speech and Language* 45. 326–347. https://doi.org/https://doi.org/10.1016/j.csl.2017.01.005.

Kuznetsova, Alexandra, Per B. Brockhoff & Rune H. B. Christensen. 2017. lmerTest package: Tests in linear mixed effects models. *Journal of Statistical Software* 82(13).

Lam, Tuan Q. & Duane G. Watson. 2010. Repetition is easy: Why repeated referents have reduced prominence. *Memory and Cognition* 38(8). 1137–1146. https://doi.org/10.3758/MC.38.8.1137.

Levy, Roger & T. Florian Jaeger. 2007. Speakers optimize information density through syntactic reduction. In Bernhard Schlökopf, John Platt & Thomas Hoffman (eds.), *Advances in Neural Information Processing Systems (NIPS)*, vol. 19, 849–856. Cambridge, MA: MIT Press. https://proceedings.neurips.cc/paper_files/paper/2006/file/c6a01432c8138d46ba39957a8250e027-Paper.pdf.

Lieberman, Philip. 1963. Some effects of semantic and grammatical context on the production and perception of speech. *Language and Speech* 6. 172–187. https://doi.org/https://doi.org/10.1177/002383096300600306.

Lindblom, Björn. 1990. Explaining phonetic variation: A sketch of the H&H theory. In W. J. Hardcastle & A. Marchal (eds.), *Speech Production and Speech Modelling*, vol. 55, 403–439. Dordrecht: Kluwer Academic Publishers. https://doi.org/https://link.springer.com/chapter/10.1007/978-94-009-2037-8_16.

Malisz, Zofia, Erika Brandt, Bernd Möbius, Yoon Mi Oh & Bistra Andreeva. 2018. Dimensions of segmental variability: Interaction of prosody and surprisal in six languages. *Frontiers in Communication* 3. 25. https://www.frontiersin.org/article/10.3389/fcomm.2018.00025/full. https://doi.org/10.3389/fcomm.2018.00025.

Meyer, David & Roger Schvaneveldt. 1971. Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology* 90. 227–34. https://doi.org/10.1037/h0031564.

Meyer, David E & Roger W Schvaneveldt. 1976. Meaning, memory structure, and mental processes: People's rapid reactions to words help reveal how stored semantic information is retrieved. *Science* 192(4234). 27–33.

Montani, Ines, Matthew Honnibal, Matthew Honnibal, Adriane Boyd, Sofie Van Landeghem & Henning Peters. 2023. *Explosion/spacy: v3.7.2: Fixes for apis and requirements*. Zenodo. https://doi.org/10.5281/ZENODO.1212303.

Nespor, Marina & Irene Vogel. 1986. *Prosodic Phonology*. Dordrecht: Foris.

Pierrehumbert, Janet B. 2001. Exemplar dynamics: Word frequency, lenition and contrast. In Joan Bybee & Paul J. Hopper (eds.), *Frequency and the Emergence of Linguistic Structure*, John Benjamins.

Pluymaekers, Mark, Mirjam Ernestus & R. Harald Baayen. 2005a. Articulatory planning is continuous and sensitive to informational redundancy. *Phonetica* 62. 146–159. https://doi.org/10.1159/000090095.

Pluymaekers, Mark, Mirjam Ernestus & R. Harald Baayen. 2005b. Lexical frequency and acoustic reduction in spoken Dutch. *Journal of the Acoustical Society of America* 118. 2561–2569. https://doi.org/10.1121/1.2011150.

R Core Team. 2022. *R: A language and environment for statistical computing*. R Foundation for Statistical Computing Vienna, Austria. https://www.R-project.org/.

Savand, Alireza, Julien Fache, Taras Labiak, Norbert Kwizera, Jannik Hoffjann, David Miró & Silvan Laube. 2024. Alir3z4/python-stop-words. https://github.com/Alir3z4/python-stop-words.

Scaltritti, Michele, David A Balota & Francesca Peressotti. 2013. Exploring the additive effects of stimulus quality and word frequency: The influence of local and list-wide prime relatedness. *Quarterly Journal of Experimental Psychology* 66(1). 91–107. https://doi.org/https://doi.org/10.1080/17470218.2012.698628.

Shannon, Claude E. 1948. A mathematical theory of communication. *Bell System Technical Journal* 27(3). 379–423. https://doi.org/10.1002/j.1538-7305.1948.tb01338.x.

Tily, Harry, Susanne Gahl, Inbal Arnon, Neal Snider, Anubha Kothari & Joan Bresnan. 2009. Syntactic probabilities affect pronunciation variation in spontaneous speech. *Language and Cognition* 1(2). 147–165. https://doi.org/https://doi.org/10.1515/LANGCOG.2009.008.

Turk, Alice. 2010. Does prosodic constituency signal relative predictability? A Smooth Signal Redundancy Hypothesis. *Laboratory Phonology* 227–262. https://doi.org/https://doi.org/10.1515/labphon.2010.012.

Turk, Alice, Satsuki Nakai & Mariko Sugahara. 2006. Acoustic segment durations in prosodic research: a practical guide. In S. Sudhoff, D. Lenertová, R. Meyer, S. Pappert, P. Augurzky, I. Mleinek, N. Richter & J. Schliesser (eds.), *Methods in Empirical Prosody Research*, 1–28. Berlin, New York: De Gruyter. https://doi.org/https://doi.org/10.1515/9783110914641.1.

Turk, Alice & Stefanie Shattuck-Hufnagel. 2007. Multiple targets of phrase-final lengthening in American English words. *Journal of Phonetics* 34(4). 445–472. https://doi.org/https://doi.org/10.1016/j.wocn.2006.12.001.

Turk, Alice & Stefanie Shattuck-Hufnagel. 2020. *Speech Timing: Implications for Theories of Phonology, Phonetics, and Speech Motor Control*. Oxford: Oxford University Press. https://doi.org/https://doi.org/10.1093/oso/9780198795421.001.0001.

Turnbull, Rory. 2017. The role of predictability in intonational variability. *Language and Speech* 60(1). 123–153. https://doi.org/https://doi.org/10.1177/0023830916647079.

Watson, Duane G., Jennifer E. Arnold & Michael K. Tanenhaus. 2008. Tic Tac TOE: Effects of predictability and importance on acoustic prominence in language production. *Cognition* 106(3). 1548–1557. https://linkinghub.elsevier.com/retrieve/pii/S0010027707001783. https://doi.org/10.1016/j.cognition.2007.06.009.

Watson, Duane G., Mara Breen & Edward Gibson. 2006. The role of syntactic obligatoriness in the production of intonational boundaries. *Journal of Experimental Psychology: Learning, Memory and Cognition* 32. 1045–1056. https://doi.org/10.1037/0278-7393.32.5.1045.

Watson, Duane G. & Edward Gibson. 2004. The relationship between intonational phrasing and syntactic structure in language production. *Language and Cognitive Processes* 19(6). 713–755. https://doi.org/https://doi.org/10.1080/01690960444000070.

Wescoat, Michael T. 2002. *On Lexical Sharing*. Ph.D. thesis, Stanford University.

Wheeldon, Linda. 2000. Generating prosodic structure. In L. Wheeldon (ed.), *Aspects of Language Production*, 249–274. Psychology Press.

Yap, Melvin J, Chi-Shing Tse & David A Balota. 2009. Individual differences in the joint effects of semantic priming and word frequency revealed by RT distributional analyses: The role of lexical integrity. *Journal of Memory and Language* 61(3). 303–325. https://doi.org/10.1016/j.jml.2009.07.001.

Zhang, Cong, Catherine Lai, Ricardo Napoleão de Souza, Alice Turk & Tina Bögel. 2023. Language redundancy effects on f0: A preliminary controlled study. In *Proceedings of the 20th International Congress of Phonetic Sciences ICPhS 2023*, Guarant International. https://guarant.cz/icphs2023/877.pdf.

Zhao, Tianyi. Forthcoming. Language redundancy effects on prosodic prominence and f0 pattern in Standard German. In *Proceedings of P&P 2024*, Halle (Saale), Germany. https://doi.org/http://dx.doi.org/10.25673/116710.

Zhao, Tianyi, Tina Bögel, Alice Turk & Ricardo Napoleão de Souza. 2024. Language redundancy effects on the prosodic word boundary strength in Standard German. In *Proceedings of Speech Prosody 2024*, 1085–1089. https://doi.org/10.21437/SpeechProsody.2024-219.

Zipf, George. 1936. *The Psychobiology of Language*. London: Routledge.

Zipf, George. 1949. *Human Behavior and the Principle of Least Effort*. New York: Addison-Wesley.