

# Vector Space Morphology with Linear Discriminative Learning

Yu-Ying Chuang, Mihi Kang, Xuefeng Luo, and R. Harald Baayen  
University of Tübingen

## 1 Introduction

Linear Discriminative Learning (LDL) is the statistical engine implementing the mappings between form and meaning in a computational model of the mental lexicon, the ‘Discriminative Lexicon’ (DL) laid out in [Baayen et al. \(2019\)](#). Given numeric representations for words’ forms as well as numeric representations for their meanings, the mathematics of multivariate multiple regression are used to transform form vectors into meaning vectors (comprehension), and to transform meaning vectors into form vectors (production). The beta coefficients of the regression models are equivalent to the weights on the connections in fully connected simple networks with input units, output units, and no hidden layers, when these network is trained to perfection ([Chuang et al., 2020a](#); [Shafaei-Bajestan et al., 2021](#)). In other words, the regression coefficients represent the best estimates (in the least squares sense) of the endstate of incremental learning. In this study, we focus on this endstate.

Although extremely simple, we have found the DL and its implementations of vector space morphology with LDL surprisingly effective for modeling a wide range of phenomena in lexical processing (see [Chuang and Baayen, 2021](#), for an overview). In what follows, we present three new case studies illustrating the potential of vector space morphology for morphological theory and the study of lexical processing. We first illustrate how LDL can be used to model Korean verb inflection. Our goal here is to show that computationally, the algorithms of LDL work for a non-trivial dataset of a language with complex morphology. We then show how LDL can be used to predict morphological priming in auditory comprehension, using data from Dutch ([Experiment 1 of Creemers, 2019](#)). Here our goal is to demonstrate that our approach offers predictive precision at the item level for data that prima facie would seem to support theories positing obligatory morphological decomposition as the first step in the lexical processing of complex words. Our third case study addresses the production of Mandarin words, focusing on acoustic duration in spontaneous speech as response variable. The goal of this final study is to illustrate how measures can be derived from LDL networks that make it possible to enhance prediction accuracy beyond what is possible with standard predictors such as word length in phones and frequency of occurrence. The general discussion provides reflection on the implications of our approach for morphological theory, theories of lexical processing, and experimental design of studies on morphological processing.

## 2 Vector space morphology for Korean verbs

Korean is generally considered to be a language isolate, possibly distantly related to Altaic languages such as Japanese. It has a complex morphology that challenges linguistic analysis in many ways, even though it is generally described to be agglutinative. The Korean verb system is especially rich. Although there is no agreement about how to best analyse Korean verbs (see [Yang, 1994](#), for detailed discussion), the seven categories (roughly following [Sohn, 2013](#)) listed in [Table 1](#) provide an overview of the factors that jointly determine a verb’s form.

Table 1: Inflectional and derivational dimensions of the Korean verb system. The values in italics pertain to example (1).

derivation	(1)	active (unmarked), causative, <i>passive</i>
subject honorifics	(2)	plain (unmarked), <i>honorific</i>
tense	(3)	present (unmarked), <i>past</i> , future
modality	(4)	neutral (unmarked), volition, <i>conjecture</i>
addressee honorifics	(5)	plain (unmarked), <i>formal</i> , polite, intimate
mode	(6)	neutral (unmarked), subjunctive, <i>retrospective</i> , ...
illocutionary force	(7)	declarative, <i>inquisitive</i> , imperative, propositive

Subject honorifics express the relative status of the speaker and the subject of the clause. If the subject of the sentence has superior status, the subject typically is an older relative, a stranger, an employer, or a teacher. Addressee honorifics honor the relative status of speaker and interlocutor, they are used to show respect towards a speaker’s or writer’s audience. Addressee honorifics are also referred to as speech levels, and we will use this terminology. Up to six speech levels have been distinguished in the literature but in Table 1 we only list the four most common ones. The category labeled ‘mode’ comprises a great many nuances that are realized by means of a wide range of exponents. The language has a four-way contrast for different kinds of illocutionary force. Example (1) presents a sentence with a verb on which all these seven inflectional categories are realized, the numbers refer to the inflectional and derivational dimensions listed in Table 1, and the values for these dimensions are presented in italics in this table.

- (1) 그 분이 잡히시었었겠습니까?  
ku pwun-i cap-hi-si-ess-ess-keyss-sup-ti-kka?  
the person-NM catch-(1)-(2)-(3-3)-(4)-(5)-(6)-(7)  
‘Did you feel (guess) that he had been caught?’ (Sohn, 2013)

Most Korean verb forms do not overtly express all seven morphological categories. When particular categories are not expressed, the ordering of the inflectional exponents for the other categories remains the same as in (1). Table 2 shows some possible combinations of the inflectional verb forms varying in speech levels (the first four examples) and illocutionary force (the other four examples at the bottom).

As mentioned above, theories vary in terms of how to decompose Korean verbs into stems and exponents (e.g., Lee, 1989; Choi, 1929; Martin, 1960; Yang, 1994). The segmentation in (1) follows Sohn (2013). The segmentation of the inflectional features and their semantic interpretation differ from scholar to scholar. For example, a verbal suffix sequence such as -(스)읍니다(-(s)upnita) is analyzed as:

1. sup-ni-ta -읍니다, formal speech level+indicative mood+declarative (Sohn, 2013);
2. upni-ta -읍니다, formal speech level+declarative (Yang, 1994);
3. up-nita -읍니다, formal speech level+declarative (Choi (1929), cited in Yang (1994));
4. up-ni-ta -읍니다, formal speech level+indicative mood+declarative (Lee, 1989).

Table 2: Examples of inflected variants of the verb 만나(manna), ‘to meet’. The top four examples vary in speech levels, the bottom four examples illustrate inflection for illocutionary force. Dots represent approximate exponent boundaries. SH: subject honorifics; SL: speech level.

manna-si.ess.up.nita	만나셨습니다	stem-honorific(SH)-past-formal(SL)-declarative
manna-si.ess.eyo	만나셨어요	stem-honorific(SH)-past-polite(SL)-declarative
manna-si.ess.e	만나셨어	stem-honorific(SH)-past-intimate(SL)-declarative
manna-si.ess.ta	만나셨다	stem-honorific(SH)-past-plain(SL)-declarative
manna-nta	만난다	stem-plain(SH)-present-plain(SL)-declarative
manna-ni	만나니	stem-plain(SH)-present-plain(SL)-inquisitive
manna-la	만나라	stem-plain(SH)-present-plain(SL)-imperative
manna-ca	만나자	stem-plain(SH)-present-plain(SL)-propositive

Although Korean is described as agglutinative, i.e., with a morphological system that uses sequences of affixal morphemes as building blocks, the above examples clarify that a strict morpheme-based analysis is not straightforward. Some exponents are clearly fusional, such as -라(-la) in 만나라 (manna-la), which realizes both plain speech level and imperative illocutionary force. Other exponents that realize both plain speech level and illocutionary force are -니(-ni), -자(-ca), and -ㄴ다 (-nta). Both stems and exponents may show allomorphy or even suppletion. Thus, the infinitive of the verb ‘to eat’ has two allomorphs, 먹(mek), and 먹어(meke). Likewise, -ㄴ다(-nta, plain speech level, declarative illocutionary force) is found after verb stems ending in vowels, but it appears as -는다(-nunta) after consonants. Stem suppletion is found for, e.g., the verb ‘to eat’, which when used with plain honorific speech requires the allomorph 먹(mek), but when used with subject honorifics is realized as 드시(tusi). Unsurprisingly, in second language teaching of Korean, learners are guided to pick up a range of low-level markers for inflectional functions, without calling explicit attention to theoretical position classes as exemplified in (1).

The non-trivial morphology of Korean raises the question of whether the simple linear mappings of discriminative learning can actually master this system, without ever having to define morphemes or exponents. To address this question, the present study focuses on verb forms realizing combinations of the four most frequently used inflectional categories: subject honorifics (labelled as ‘honorifics’ in our dataset), tense, speech level, and illocutionary force. The feature values that we implemented are those listed above in Table 1. Our dataset also includes some derived words realizing passives or causatives, as these features are commonly used in daily speech. The combinations of these features result in 59 inflected forms for a standard verb. Semantically impossible combinations are not included (e.g., past propositive or future imperative). In this dataset, which in total comprises 27,258 word forms of in all 462 verbs, words’ forms are represented by syllables. These syllables do not necessarily coincide with Hangul characters. Although Hangul characters are designed to represent syllables, due to assimilation, contraction, and resyllabification, the spoken forms of words can diverge considerably from the Hangul orthography. We opted for phonological transcriptions that follow the spoken forms. Table 3 presents four rows of our dataset, which provides the Hangul written form, a phonemic representation with underscores indicating syllable boundaries, an identifier for the verb lexeme, and the feature specifications for honorifics, tense, speech level, and illocutionary force. In our dataset, differences in intonation, which can distinguish between words that are otherwise homophones but that have different illocutionary force (e.g., 먹어요 (mekeyo), ‘to eat’, present tense, polite speech level, declarative/inquisitive/propositive/imperative), are not

Table 3: Four Korean verb forms and their phonological and morphological properties that form the input for modeling with LDL.

	Hangul	Word	Lexeme	Honorifics	Tense	SpeechLevel	IllocutionaryForce
1	고릅니다	go_rUm_ni_da	gorUda	plain	present	for	dec
2	고릅니까	go_rUm_ni_kka	gorUda	plain	present	for	inq
3	고르십시오	go_rU_sip-_syo	gorUda	hon	present	for	imp
4	고릅시다	go_rUp-_si_da	gorUda	plain	present	for	pro

yet indicated.

In what follows, we will introduce the main concepts behind LDL step by step, without going into mathematical detail, but providing the code for setting up and running the model using the **JudiLing** package for the Julia language (Luo et al., 2021).<sup>1</sup> The choice for Julia as programming language is motivated by the fact that this language is extremely well optimized for numerical calculation. As a consequence, building and evaluation with **JudiLing** is much faster than is possible with the **WpmWithLdl** package for R.<sup>2</sup>

Before starting the actual modeling, after having started up Julia, we first load the **JudiLing** package and the packages for handling csv files and data frames. The following code snippet assumes that our dataset, `korean.csv`, is available in the current working directory.

```
using JudiLing, CSV, DataFrames
korean = DataFrame(CSV.File("korean.csv"));
```

In order to set up mappings between words' forms and their meanings, we need to define numeric representations for form and meaning. We first consider setting up a matrix with row vectors that represent words' forms. There are many ways in which form vectors can be defined. In the present study, we set up form vectors that specify which pairs of consecutive syllables occur in a word. That is, we build a form matrix  $C$  with as many rows as there are words, and with as many columns as there are syllable pairs. We indicate the presence of a syllable pair with 1, and its absence with 0. Constructing this matrix with **JudiLing**'s `make_cue_matrix` function is straightforward.

```
C_obj = JudiLing.make_cue_matrix(korean, grams=2, target_col=:Word, tokenized=true,
                                sep_token="_", keep_sep=true);
```

The directive `grams=2` tells the model to construct di-syllables, the directive `target_col` specifies in which column of the dataset word form information can be found, and the `sep_token` directive specifies that the underscore represents the syllable boundary marker. The object returned by `make_cue_matrix`, `C_obj`, contains several useful data structures, of which the  $C$  matrix (`C_obj.C`) is of interest to us here.

```
JudiLing.display_matrix(korean, :Word, C_obj, C_obj.C, :C, nrow=4, ncol=9)
```

		#_go	go_rUm	rUm_ni	ni_da	da_#	ni_kka	kka_#	go_rU	rU_sip-
1	go_rUm_ni_da	1	1	1	1	1	0	0	0	0
2	go_rUm_ni_kka	1	1	1	0	0	1	1	0	0
3	go_rU_sip-_syo	1	0	0	0	0	0	0	1	1
4	go_rUp-_si_da	1	0	0	0	1	0	0	0	0

<sup>1</sup>For installation of the package and the packages that it depends on, see the on-line manual at <https://megamindhenry.github.io/JudiLing.jl/stable/#Installation>.

<sup>2</sup>The latest version of this package (`WpmWithLdl_1.4.6.tar.gz`) is available at <https://osf.io/xq92s>.




Figure 1: Vector addition shifts the position of a vector in semantic space. In this example, the vectors of the two verbs in present tense forms (*gam\_ni\_da* ‘to go’, *go\_rUm\_ni\_da* ‘to pick’) are added with the vector of PAST, which moves them from the PRESENT area to the PAST area.

The second modeling step is to construct the matrix  $\mathbf{S}$  that represents words’ meanings in numeric form, as in distributional semantics (Landauer and Dumais, 1997; Mikolov et al., 2013; Mitchell and Lapata, 2008). These vectors can be extracted from corpora using standard methods from computational linguistics. In this study, however, we simulate these vectors, but we make sure that they properly represent inflectional contrasts (see Baayen et al., 2018; Chuang et al., 2020b, for further examples). This is accomplished by summing the simulated semantic vectors of the verb stem and the simulated vectors representing the various inflectional features realized in a given inflected form. Thus, the semantic vector of the first word *go\_rUm\_ni\_da* is obtained as follows:

$$\overrightarrow{go\_rUm\_ni\_da} = \overrightarrow{gorUda} + \overrightarrow{plain} + \overrightarrow{present} + \overrightarrow{for} + \overrightarrow{dec}. \quad (1)$$

All elementary simulated semantic vectors consist of random numbers following Gaussian distributions. By adding inflectional vectors to base vectors, the meanings of base words are shifted systematically, and in different ways that are specific to the inflectional meanings that are realized in conjunction with that base word. Figure 1 illustrates this concept for tense inflection. By adding the vector of PAST to the vectors of the verbs in the present tense, we move their positions from the area in space where present tense forms are located to the area in space where past-tense forms cluster.

Simulating morphologically-aware semantic vectors with **JudiLing** is straightforward. All we need to do is specify the dataset, the name of the column in the dataset that provides the verb stems (lexemes), and the names of the columns that specify inflectional feature values.

```
S = JudiLing.make_S_matrix(korean,
                           [ :Lexeme ],
                           [ :Honorifics, :Tense, :SpeechLevel, :IllocutionaryForce ],
                           ncol=size(C_obj.C, 2));
```

The `ncol` directive specifies the number of columns of  $\mathbf{C}$ , ensuring that the length of the form vectors and the semantic vectors are identical. The upper-left corner of  $\mathbf{S}$  is displayed as follows:

```
JudiLing.display_matrix(korean, :Word, C_obj, S, :S, nrow=4, ncol=7)
```

		S1	S2	S3	S4	S5	S6	S7
1	go_rUm_ni_da	-18.5863	6.61192	4.24111	2.01868	9.9817	14.7086	11.1827
2	go_rUm_ni_kka	-21.4642	-2.87628	3.51315	4.98055	11.0686	15.6375	1.0228
3	go_rU_sip_syo	-30.2516	-2.66979	2.68477	1.4475	11.2798	6.68343	8.08859
4	go_rUp_si_da	-23.2687	-4.55982	0.965755	5.25853	12.1366	14.1624	1.43157

With  $\mathbf{C}$  and  $\mathbf{S}$  in hand, we can now proceed to set up the mappings from form to meaning (comprehension), and from meaning to form (production). We first compute the comprehension network  $\mathbf{F}$ , which predicts word meanings from word forms, by solving  $\mathbf{S} = \mathbf{CF}$ . **JudiLing**'s function `make_transform_matrix` estimates  $\mathbf{F}$  in a numerically highly optimized way.

```
F = JudiLing.make_transform_matrix(C_obj.C, S);
```

The  $\mathbf{F}$  matrix specifies the connection weights (in network parlance) or equivalently the beta weights (in the parlance of multivariate multiple regression) of all the di-syllable cues (rows) to all semantic dimensions (columns). In other words,  $\mathbf{F}$  is a trained network, or equivalently, a multivariate multiple regression model, that takes a verb's form vector  $\mathbf{c}$  as input and predicts its semantic vector  $\hat{\mathbf{s}}$ . Restating this for all words jointly, we obtain the matrix with predicted semantic vectors  $\hat{\mathbf{S}}$  by multiplying  $\mathbf{C}$  with  $\mathbf{F}$ :

```
Shat = C_obj.C * F;
```

How successful is this mapping from form to meaning? How close are the predicted semantic vectors to the gold standard vectors in  $\mathbf{S}$ ? We can evaluate closeness by means of the Pearson correlation, and evaluate the mapping to be successful for a given word  $i$  if its predicted semantic vector  $\hat{\mathbf{s}}_i$  has the highest correlation with the gold standard vector  $\mathbf{s}_i$ , that is, when

$$r(\hat{\mathbf{s}}_i, \mathbf{s}_i) = \max_j \{r(\hat{\mathbf{s}}_j, \mathbf{s}_i)\}.$$

The `eval_SC` function of **JudiLing** evaluates model accuracy in this way.

```
JudiLing.eval_SC(Shat, S, korean, :Word)
0.9918
```

For the present dataset, with over 27 thousand word forms, we obtain an accuracy of 99%.

For modeling the production of Korean verb forms, we first derive the weight matrix  $\mathbf{G}$ , associated with the network that maps meanings  $\mathbf{S}$  to forms  $\mathbf{C}$ . With  $\mathbf{G}$  in hand, we calculate the predicted form matrix  $\hat{\mathbf{C}}$ .

```
G = JudiLing.make_transform_matrix(S, C_obj.C);
Chat = S * G;
```

The  $\hat{\mathbf{C}}$  matrix specifies, for each word (rows) the amount of estimated semantic support that the di-syllables receive for that word. In the following tabulation, the di-syllable `go_rUm` is well-supported for the first and second words, but not for the third and fourth ones.

```
JudiLing.display_matrix(korean, :Word, C_obj, Chat, :Chat, nrow=4, ncol=7)
```

		#_go	go_rUm	rUm_ni	ni_da	da_#	ni_kka	kka_#
1	go_rUm_ni_da	0.747749	0.184328	0.188502	0.656815	0.757723	0.190278	0.190721
2	go_rUm_ni_kka	0.753465	0.185616	0.186941	0.211225	0.354903	0.541299	0.535272
3	go_rU_sip_syo	0.753545	0.027506	0.012114	0.225991	0.324489	0.223644	0.220002
4	go_rUp_si_da	0.753558	0.027182	0.031723	0.206195	0.728115	0.184292	0.185375

We now know which di-syllables are most likely to be part of a given word, but at this point, we have no criterion for selecting those di-syllables that actually occur in a given word, and we also have no information about their proper order. Since di-syllables contain partial information about syllable order, we can weave syllables together. There are several ways in which this can be accomplished. Here, we make use of an algorithm that first learns to predict, for a given word, and for all possible positions, which di-syllable is most appropriate for that position. Given the top candidates, the algorithm then weaves together di-syllables into words, which typically results in several possible word candidates. The algorithm then selects the candidate that, when presented to the comprehension network, best approximates the targeted meaning (Baayen et al., 2019, refer to this as synthesis by analysis). The `learn_paths` function implements this algorithm.

```
res = JudiLing.learn_paths(korean, C_obj, S, F, Chat);
```

We evaluate production accuracy by comparing the predicted forms with the targeted forms. Production accuracy is 90%.

```
JudiLing.eval_acc(res, C_obj)
0.9006
```

What we have shown, for a non-trivial Korean dataset of complex inflected verbs, is that the simple mappings of the Discriminative Lexicon model work with surprisingly high levels of accuracy for both production and for comprehension. It is beyond the scope of the present study to evaluate model accuracy on held-out data, for a first exploration of the productivity of LDL networks, see Chuang et al. (2020b). Now that we have shown that LDL can indeed learn the mappings between form and meaning without requiring stems, morphemes, or exponents, the question arises of whether the model can generate useful predictions for experimental data on lexical processing. The next section takes on this question for comprehension.

### 3 Comprehension

Creemers et al. (2020) used the primed auditory lexical decision task to assess processing costs for Dutch complex verbs with initial particles or prefixes. This study observed equivalent priming effects for semantically transparent and semantically opaque verbs, as compared to control conditions, replicating earlier work on German verbs (Smolka et al., 2007; Smolka and Eulitz, 2018). These results dovetail well with theories positing that morphological processing involves automatic decomposition of forms into their constituents (Taft, 2004; Stockall and Marantz, 2006; Lewis et al., 2011; Marantz, 2013), irrespective of their semantics.

LDL does not decompose visual or auditory input into morphological constituents. Instead, it decomposes its input into finer-grained, overlapping sublexical units such as trigrams, triphones, or syllables (for modeling audio input with LDL, see Shafaei-Bajestan et al., 2021). Since the model sets up simple mappings between sublexical units and semantic vectors, it is a priori not self-evident that it can predict priming effects that suggest the presence of form effects in the absence of semantic effects. Baayen and Smolka (2020) addressed this question for German, and showed, using the Naïve Discriminative Learning (NDL) model of Baayen et al. (2011), that priming effects are inversely proportional to the extent that the trigrams in a prime word pre-activate the semantic unit representing the meaning of the target word. The downside of the NDL model is that all word meanings are construed as orthogonal, thanks to its one-hot encoding of words' meanings, and thus

completely unrelated semantically. Hence, it might be argued that semantics does not really come into its own in this model, and that the NDL model of Baayen and Smolka (2020) is simply capturing pure form effects. In what follows, we therefore move from NDL to LDL, and apply the latter model to the Dutch complex verbs used in Experiment 1 of Creemers et al. (2020). Importantly, unlike for our Korean dataset, where we simulated semantic vectors, in this modeling study we made use of empirical semantic vectors. If LDL correctly predicts the observed priming effects, then this cannot be due to the model not having access to proper corpus-based word embeddings.

Input for modeling was a set of 7803 Dutch words, selected from the lemma database in CELEX (Baayen et al., 1995). For a word to be included, it had to have a frequency of occurrence exceeding 100 (per 42 million) and have at most two constituents, or it had to be listed as a prime or target in Experiment 1 of Creemers et al. (2020). The resulting dataset is loaded into Julia as follows:

```
dutch = DataFrame(CSV.File("dutch.csv"));
```

For form vectors, we opted for using phone trigrams as sublexical units of form.

```
C_obj = JudiLing.make_cue_matrix(dutch, grams=3, target_col=:phon);
size(C_obj.C)
(7803, 6356)
```

For the semantic matrix  $S$ , we used the 300-dimensional vectors constructed with `fasttext` as available at <https://fasttext.cc/docs/en/crawl-vectors.html>. The file `S_dutch.txt` contains the semantic vectors for the words in our dataset, in exactly the order in which they occur in `dutch`. The first column of `S_dutch.txt` specifies the word, and the remaining 300 columns specify the word embedding. The vectors and the corresponding words are loaded into Julia as follows:

```
S, words = JudiLing.load_S_matrix("S_dutch.txt")
```

$S$  is then the semantic matrix, and the object `words` lists the words in the dataset. After obtaining the transformation matrix  $F$  and the predicted semantic vectors  $\hat{S}$ ,

```
F = JudiLing.make_transform_matrix(C_obj.C, S);
Shat = C_obj.C * F;
```

we calculated all pairwise correlations of gold standard and predicted semantic vectors, and observed that for 91% of words, the predicted semantic vector was the one that was closest to its targeted gold standard. The following code snippet calculates both the accuracy and the matrix of pairwise correlations  $R$  of words' estimated and gold standard semantic vectors, which we will need below.

```
accuracy, R = JudiLing.eval_SC(Shat, S, dutch, :phon, R=true);
accuracy
0.909
```

In order to model lexical priming, we conceptualize the reader's or listener's current semantic state as a pointer to a location in semantic space. An experimental trial can then be seen as beginning with a semantic pointer at the origin of this space. Presentation of the prime moves this pointer to the position in semantic space where the prime's semantic vector is approximately located. We may expect that when this position is closer to that of the target word, the target word is primed more strongly. We evaluated closeness with the correlation measure. For all prime-target pairs, we extracted from  $R$  the correlation  $r_{ij}$  of the predicted vector for prime  $i$  and the






Figure 2: Distributions of empirical reaction times (left), prime-to-target approximations (center) and prime-to-target approximations with  $\mathcal{N}(0, 0.3)$  response noise added.

gold standard vector for target  $j$ . In what follows, we refer to these correlations as prime-to-target approximations (PTA).

Table 4 and Figure 2 show that these correlations capture the priming effects reported by Creemers et al. (2020). As expected, greater values of PTA predict shorter RTs. The correlation between PTA and mean by-item log-transformed reaction time (as made available by Creemers, 2019) was  $-0.3$  ( $t(141) = -3.67, p = 0.0003$ ). It is clear that LDL succeeds in modeling morphological priming data, replicating Baayen and Smolka (2020) (see also Milin et al., 2017). It does so without ever having to parse forms into disjunct sets of morphological constituents.

Table 4: Predicting prime-target approximation from prime type. The reference level for prime type is morphologically related but semantically unrelated prime (m); c: control primes, ph: phonologically related primes, ms: morphologically and semantically related primes.

	Estimate	Std. Error	t value	Pr(> t )
Intercept	0.5533	0.0166	33.40	0.0000
prime_type = c	-0.2464	0.0236	-10.44	0.0000
prime_type = ms	0.0749	0.0234	3.20	0.0017
prime_type = ph	-0.2187	0.0234	-9.34	0.0000

Having illustrated how measures derived from LDL can be used to predict lexical processing costs as gauged by a priming paradigm, in the next section we consider how the model can be used to predict words’ acoustic durations.

## 4 Speech production: Acoustic duration for Taiwan Mandarin words

One measure which is informing theories of speech production is words’ acoustic durations (Gahl, 2008; Gahl et al., 2012; Plag et al., 2017; Tomaschek et al., 2019). In this study, we investigated the

acoustic durations of words in Taiwan Mandarin. From a corpus of spontaneous Taiwan Mandarin speech (Fon, 2004), we selected all 7,349 one- to five-syllable words for which `fasttext` semantic vectors are available. In the following snippet of Julia code, we assume the dataset `mandarin.csv` is available in the working directory.

```
mandarin = DataFrame(CSV.File("mandarin.csv"));
```

An overview of this dataset is presented in Table 5.

Table 5: The first four rows of the Mandarin dataset. Glosses for the words from top to bottom: ‘one’, ‘one by one’, ‘a bit’, and ‘abruptly’.

	word	pinyin	phones
1	一	yi1	i1
2	一一	yi1_yi1	i1.i1
3	一下	yi1_xia4	i1.x.ia4
4	一下子	yi1_xia4_zi5	i1.x.ia4.z.ii5

In this dataset, the second column (`pinyin`) provides the transcription of the words according to the standard romanization system. For modeling, however, we used the more detailed transcription of phones in the third column (`phones`). Here, phones are separated by periods, and Mandarin lexical tones are indicated by numbers next to the vowels. As in the previous study on Dutch, we used triphones to build form vectors.

```
C_obj = JudiLing.make_cue_matrix(mandarin, grams=3, target_col=:phones,
                                tokenized=true, sep_token=".", keep_sep=true);
```

Given a file with `fasttext` vectors (`S_mandarin.txt`), with words’ vectors listed in the same order as the words in the dataset of `mandarin`, we load the semantic matrix into Julia,

```
S, words = JudiLing.load_S_matrix("S_mandarin.txt")
```

and then calculate the transformation matrix  $F$ , the matrix with predicted semantic vectors  $\hat{S}$ :

```
F = JudiLing.make_transform_matrix(C_obj.C, S);
Shat = C_obj.C * F;
```

Comprehension accuracy is high:

```
JudiLing.eval_SC(Shat, S, mandarin, :phones)
0.9877
```

For modeling production, we first estimate the mapping  $G$  and the matrix with predicted triphone supports  $\hat{C}$ ,

```
G = JudiLing.make_transform_matrix(S, C_obj.C);
Chat = S * G;
```

and then use the `learn_paths` function to weave the triphones into words.

```
res, gpi = JudiLing.learn_paths(mandarin, C_obj, S, F, Chat, threshold=0.01,
                               Shat_val=Shat, check_gold_path=true);
```

The directive `check_gold_path`, when set to `true`, requests the function to keep track of the amount of support that the triphones of a given word receive in different positions, and outputs a second object `gpi`. As will become clearer later, we will need this object to derive network measures. We note here that for production, accuracy is also satisfactory:

```
JudiLing.eval_acc(res, C_obj)
0.9291
```

From this model of the Mandarin mental lexicon, we derived two novel measures that are predictive for Mandarin words’ acoustic durations.

The first measure is derived from the production network. Recall that the weight matrix of this network,  $\mathbf{G}$ , has semantic dimensions on rows and triphones on the columns. The column vectors of  $\mathbf{G}$  quantify the extent to which triphones contribute to realizing the meanings of their carrier words. To make this more concrete, consider an example lexicon with just four English words: *tree*, *bee*, and their plural forms *trees*, *bees*. Model setup is presented in Table 6: for cues we used diphones, and their semantic vectors, using simulated vectors, are the sum of their respective base word vector (TREE/BEE) and number vector (SG/PL). Following the same procedure as described above, we obtained the  $\mathbf{G}$  matrix, the production network mapping meanings onto forms. Table 7 presents the column vectors of  $\mathbf{G}$ , once for *trees* (left) and once for *bees* (right), together with a vector of zeroes representing the origin.

Table 6: Example lexicon with four words. The second and third columns specify words’ diphones and how their semantic vectors are constructed.

Word	diphones	semantics
tree	#t tr ri i#	$\overrightarrow{\text{TREE}} + \overrightarrow{\text{SG}}$
bee	#b bi i#	$\overrightarrow{\text{BEE}} + \overrightarrow{\text{SG}}$
trees	#t tr ri iz z#	$\overrightarrow{\text{TREE}} + \overrightarrow{\text{PL}}$
bees	#b bi iz z#	$\overrightarrow{\text{BEE}} + \overrightarrow{\text{PL}}$

From a neurobiological perspective, we can regard the  $\mathbf{G}$  matrix, in the roughest sense, as a production neural network, with semantic features and diphone cues as representational units of neurons (or collections of neurons) in the cortex. As neurons have a location in the cortex, we therefore also examine what a sensible topological organization of the diphone units of our network would be. Note that the order in which diphones are listed in the columns of the  $\mathbf{G}$  matrix is completely arbitrary — all that counts for the mathematics is their identity and their connectivity. By contrast, neurons in the cortex have non-arbitrary topologies (see, e.g., [Ferro et al., 2011](#)). There are many different ways of projecting representations in a high-dimensional space to a two-dimensional surface that could serve as a rough approximation of a (highly simplified) cortical surface (see also [Heitmeier and Baayen, 2020](#); [Baayen et al., 2018](#)). Here we approximated such a surface by applying principal component analysis to the column vectors of  $\mathbf{G}$ . As can be seen from Figure 3, the stem diphones of *tree* (#t, tr, ri) appear together on the left, marked by green circle, whereas the stem diphones of *bee* (#b, bi) cluster in the orange circle on the right. Interestingly, the diphone i# on the one hand, and those of iz and z# on the other, also are located in two different positions, at the top and bottom respectively.

To understand why the diphones cluster in this way, we have to go back the column vectors of  $\mathbf{G}$  (henceforth diphone vectors). Since the diphone vectors share the same dimensions as the semantic

Table 7: Cue weights in the production network  $\mathbf{G}$ . Table (a) and (b) present diphones required to produce *trees* and *bee*, respectively. The bottom row indicates pair-wise Euclidean distances between diphone vectors, starting from the origin (all-zero vector) to the last diphone of the word.

(a) <i>trees</i>							(b) <i>bee</i>			
	origin	#t	tr	ri	iz	z#	origin	#b	bi	i#
S1	0	-0.01	-0.01	-0.01	-0.03	-0.03	0	-0.01	-0.01	0.01
S2	0	-0.06	-0.06	-0.06	-0.11	-0.11	0	0.04	0.04	0.08
S3	0	0.05	0.05	0.05	-0.04	-0.04	0	-0.11	-0.11	-0.02
S4	0	0.07	0.07	0.07	0.03	0.03	0	-0.03	-0.03	0.00
S5	0	-0.02	-0.02	-0.02	0.01	0.01	0	0.07	0.07	0.03
S6	0	0.00	0.00	0.00	-0.02	-0.02	0	-0.02	-0.02	0.00
S7	0	-0.07	-0.07	-0.07	0.00	0.00	0	0.04	0.04	-0.03
S8	0	-0.05	-0.05	-0.05	-0.01	-0.01	0	0.00	0.00	-0.05
distance	0.137	0	0	0.143	0		0.147	0	0.143	




Figure 3: Projection of diphone vectors onto a two-dimensional ‘cortical’ plane, using principal component analysis. For *trees*, the total Distance Travelled is  $d1 + 0 + 0 + d2 + 0$ . For *bee*, the distance is  $d3 + 0 + d4$  (see also Table 7).

vectors, we can examine how they are semantically related by calculating the pair-wise correlations of the diphone vectors and the vectors of all base words and inflectional meanings — TREE, BEE, SG, PL, referred to as lexomes in Baayen et al. (2019). The resulting correlation matrix is presented in Table 8. The diphones of *tree* are best correlated with  $\overrightarrow{\text{TREE}}$ , and the diphones of *bee* are best correlated with  $\overrightarrow{\text{BEE}}$ . Furthermore, *i#* is strongly correlated with  $\overrightarrow{\text{SG}}$ , whereas *iz* and *z#* correlate well with  $\overrightarrow{\text{PL}}$ . The correlations in bold in Table 8 can therefore be understood as measuring the functional load of the diphones. In this example, *i#* realizes primarily singular number, *z#* realizes

Table 8: Correlations of the diphone vectors in  $\mathbf{G}$  and the semantic vectors of the base word and number lexomes. The highest correlation in each row is highlighted in bold.

	TREE	BEE	SG	PL
#t tr ri	<b>0.55</b>	-0.44	0.26	0.13
#b bi	-0.42	<b>0.72</b>	0.15	0.28
i#	0.04	0.19	<b>0.85</b>	-0.17
iz z#	0.11	0.18	-0.29	<b>0.68</b>

plural number, and #b and bi realize  $\overrightarrow{\text{BEE}}$ .<sup>3</sup> It is surprising how well the network succeeds in figuring out, for this agglutinative example, which bits of form have what purpose, without ever having to tell the model explicitly about stems and exponents.<sup>4</sup>

Is lexical processing, specifically durations of word production in this study, influenced by the proximity of diphones in such a hypothesized cortical 2-D plane? We reasoned that when diphones have very similar functional load, the costs of evaluating their contribution to production is smaller than when diphones are at greater distance of each other. This hypothesis leads to our first measure for predicting acoustic durations, the total distance travelled in semantic space when tracing a word’s successive diphones (henceforth **Distance Travelled**). Consider again Table 7(a), which arranges the diphones in the order in which they drive articulation. Assuming that we start from the null meaning (represented by the origin, an all-zero vector), we can calculate the Euclidean distance from one successive diphone vector to the next, as indicated in the last row of the table. From the origin to the vector of #t, this distance is 0.137. In Figure 3, this distance is labeled as **d1**, in green. From #t to tr, and from tr to ri, the distance is zero: these three diphones are associated with the TREE meaning to the same extent, and have the same location in Figure 3. Hence for these diphones, no further distance penalty is incurred. However, for the next transition, as iz is located further down in the hypothesized cortical plane, we obtain a non-zero distance (**d2** in Figure 3). Finally, z# is in the same position as iz, hence the distance between the last two diphones is zero. The total Distance Travelled for *trees* is 0.137 (d1) + 0 + 0 + 0.143 (d2) + 0 = 0.28. For *bee* (Table 7(b)), the total Distance Travelled is 0.147 (d3) + 0 + 0.143 (d4) = 0.29. (For more realistic datasets, distances between diphones or triphones will rarely be zero, as the network weights for transitions represent compromises between many conflicting constraints.) The larger the value of Distance Travelled, the more heterogeneous the diphone or triphone vectors are, the more variegated the meanings are that they have to realize, and the lower the within-word semantic coherence is. These considerations lead us to expect that acoustic duration and Distance Travelled are positively correlated: more complex messages require longer codes (cf. Kuperman et al., 2009).

The second measure that we considered for predicting acoustic durations is much simpler. As

<sup>3</sup>These correlations clarify how those diphones are selected for production that properly realize words’ semantics. When we calculate the dot product of a row vector  $\mathbf{s}_i$  of  $\mathbf{S}$  (the sum of two lexome vectors, e.g.,  $\overrightarrow{\text{TREE}} + \overrightarrow{\text{PL}}$ ) and a diphone column vector  $\mathbf{g}_j$  of  $\mathbf{G}$ , we multiply their elements pairwise and then sum. When the two vectors are uncorrelated, this sum will be close to zero. However, when the two vectors are correlated, the dot product will tend to produce larger values. As a consequence, correlated vectors ensure that the support for a diphone in a word’s predicted form vector  $\hat{\mathbf{c}}$  will be greater.

<sup>4</sup>For the present agglutinative example, we see clusters of diphones for the two stems and the plural exponent, but the reader should be warned that for realistically sized datasets, and especially for languages with fusional morphology, such crisp and clear divisions of labor are not to be expected. Nevertheless, some albeit more diffuse clustering may still be visible.

illustrated in Section 2, the learning algorithm that orders triphones for articulation (implemented in the `learn_path` function) predicts, for a given word and a specific triphone position (ranging from 1 to the length of the longest word in the dataset), how well the different possible triphones are supported for that position. For example, for the Mandarin word 土地 (tǔdì) ‘land’, the positional triphone support for the first triphone #tu is 0.04 for position 1, and that for the second triphone tud is 0.03 for position 2. Our second measure, **Triphone Support**, is the sum of the positional supports for a words’ diphones. The larger the value of Triphone Support, the more confident the model is in terms of its prediction. We therefore expect that Triphone Support and acoustic duration enter into a positive correlation.

These two measures can be straightforwardly obtained with the **JudiLing** package:

```
distances = JudiLing.get_total_distance(C_obj, G, :G)
supports = JudiLing.get_total_support(gpi)
```

Both functions return vectors with the pertinent values for each of the words in the dataset.

Are these two measures indeed predictive for words’ acoustic durations? We first fitted a generalized additive model (GAM) (Wood, 2017) to the acoustic durations with word frequency and word length (number of phones) as predictors. Both frequency and length are strong predictors, unsurprisingly. The effect of word length is rather trivial, as all we are doing is regressing duration in milliseconds on duration in phones. The reason we include length is to ensure that our model-based measures cannot simply be reduced to word length. In other words, this model serves as our baseline model.

When Distance Travelled and Triphone Support are added as predictors to the baseline model, goodness of fit improved substantially by no less than 195 AIC units. This clarifies that the distance and support measures capture properties in production that cannot be fully accounted for by word length and frequency. Unfortunately, due to substantial collinearity, the resulting model suffers from suppression and enhancement, rendering interpretation impossible. Both our model-based measures are highly correlated with word length: for Distance Travelled,  $r = 0.81$ ,  $p < 0.0001$ , and for Triphone Support,  $r = 0.37$ ,  $p < 0.0001$ . In order to obtain a model that is better interpretable, we first regressed word length on the distance and support measures. This was achieved by fitting another GAM to word length with the two LDL measures as predictors. Instead of a standard GAM, we used a Gaussian location-scale model, which models both mean and variance of word length at the same time. The two measures are predictive for both the mean and the variance of length (see Figure 4). In general, longer words have longer distances travelled in semantic space, as well as larger support received by triphones (left two panels). Interestingly, the variance of word length also varies with the two LDL measures. While variability in length increases with Distance Travelled, it decreases with Triphone Support.<sup>5</sup> The R-squared of this model was 49%. The correlation of the residuals of this model, henceforth **Residualized Length**, with the original word length measure was 0.72.

Whereas the positive correlations of the two measures with mean length are unsurprising, their strong effect on the variance in word length is novel and of potential theoretical interest. For Triphone Support, greater support offers more precise prediction of word length (as variance in word length decreases). In other words, when the model struggles with learning triphone positions, resulting in low values, predictions for word length are more variable. By contrast, for larger values of Distance Travelled, predictions for word length apparently become less precise.

---

<sup>5</sup>Interestingly, once the two LDL measures are in the model, frequency of occurrence is no longer predictive for




Figure 4: The effects of Distance Travelled and Triphone Support on word length (number of phones). The left two panels show that mean word length increases with both predictors. The right two panels show that the variability in word length increases with Distance Travelled, but decreases with Triphone Support.

We then re-fitted the GAM, replacing word length with Residualized Length. The mid and right panels of Figure 5 plot the effects of frequency and length. Words with higher frequency have shorter durations, and phonologically longer words are produced with longer durations, as expected. What is more interesting is the interaction of the two LDL measures, as presented in the left panel. We observed a strong effect of Distance Travelled: word durations become longer with increasing Distance Travelled. This suggests that when a word’s triphones are far apart in semantic space, and possibly far apart in our hypothetical cortical plane, in which case they jointly realize more complex semantics, then the speed with which this word is articulated is slowed, reflecting greater processing costs. With regards to Triphone Support, there is a cross-over effect. For small values of Distance Travelled, the larger the support, the longer the durations. For large values of Distance

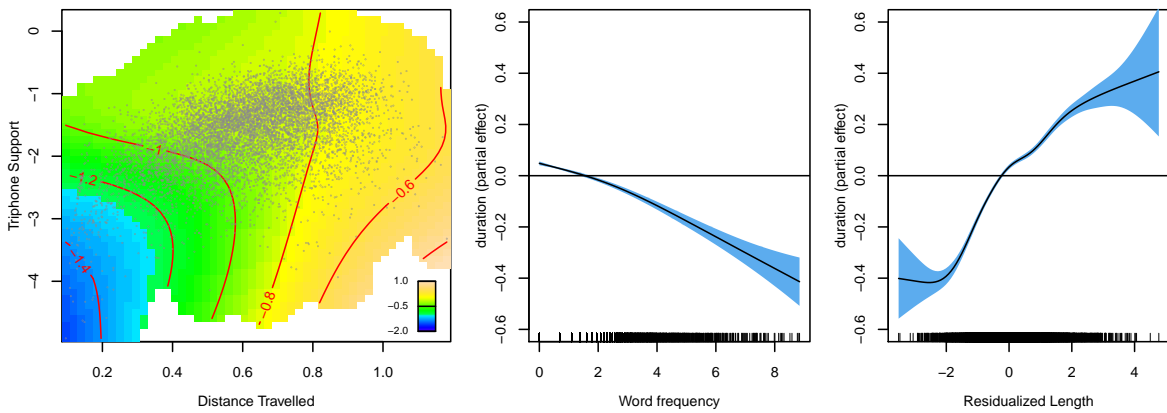


Figure 5: The interaction effect of Distance Travelled and Triphone Support on mean acoustic durations of Mandarin words produced in spontaneous speech (left panel). Warmer colors represent longer durations, and dots indicate the distribution of data points. The effects of word frequency, and Residualized Length are presented in the middle and right panels.

Travelled, however, the trend is reversed, here durations become shorter with increasing Triphone Support.

In summary, with empirical semantic vectors, LDL performed well for both comprehending and producing Mandarin words. Furthermore, when pitted against behavioral data, the measures derived from the networks are predictive for the acoustic durations of these words as actually pronounced in spontaneous speech. The results suggest that the LDL framework, despite making use of very simple networks, is useful for studying lexical processing.

## 5 General Discussion

The three case studies that we have reported illustrate why we find it fruitful to rethink morphology in terms of mappings between form and meaning spaces, using the toolkit of linear algebra and regression modeling. Although firmly rooted in Word and Paradigm Morphology, our analysis of how triphones may cluster in a two-dimensional approximation of a cortical plane suggests that, when a morphological system is agglutinative, clusters of triphones start to look very much like stems and exponents. If this finding generalizes from a toy example to realistic datasets, it would provide a bridge between Word and Paradigm Morphology and decompositional theories that are supported by neuroimaging studies suggesting that constituents are subserved by localized cell-assemblies in the brain (Bozic et al., 2007; Lewis et al., 2011; Chersi et al., 2014). At the same time, it is clear that for claims about the neurobiological reality of classical discrete morpheme and stem representations to be convincing, it will be necessary to show that there is no confound between the classical units and sublexical units such as triphones. Likewise, when designing priming studies to support the hypothesis of obligatory automatic morphological decomposition, it will be necessary to control for the kind of low-level variables that this study has shown to co-determine lexical processing, and that patterns of priming do not already emerge straightforwardly from a language’s distributional properties.

Breiman et al. (2001) famously made a distinction between two very different cultures for data analysis, contrasting statistical modeling with machine learning. Whereas machine learning is optimized for obtaining precise predictions that can be profitably put to use in industry, statistical modeling is concerned with understanding the data, and to this end builds models that could have generated the data. LDL is closer in spirit to statistical modeling than to machine learning. It capitalizes on simplicity, in the hope that this will further clarity of understanding. Undoubtedly, this simplicity comes at a cost. Although many mappings in morphology thus far have appeared to be well approximated by linear mappings, nonlinear mappings are likely to offer further precision. Interestingly, exactly the words for which LDL lacks precision are the words that are likely to be more difficult to learn also for human learners, and that hence will have greater lexical processing costs in the mental lexicon. George Box is well known for stating that all statistical models are wrong, but some are useful (Box, 1976). In the same way, even though vector space morphology as implemented with LDL requires many simplifying assumptions, it nevertheless provides a useful tool for probing the mental lexicon.

## References

Baayen, R. H., Chuang, Y.-Y., and Blevins, J. P. (2018). Inflectional morphology with linear mappings. *The Mental Lexicon*, 13(2):232–270.



- Baayen, R. H., Chuang, Y.-Y., Shafaei-Bajestan, E., and Blevins, J. (2019). The discriminative lexicon: A unified computational model for the lexicon and lexical processing in comprehension and production grounded not in (de)composition but in linear discriminative learning. *Complexity*.
- Baayen, R. H., Milin, P., Filipović Durdević, D., Hendrix, P., and Marelli, M. (2011). An amorphous model for morphological processing in visual comprehension based on naive discriminative learning. *Psychological Review*, 118:438–482.
- Baayen, R. H., Piepenbrock, R., and Gulikers, L. (1995). *The CELEX lexical database (CD-ROM)*. Linguistic Data Consortium, University of Pennsylvania, Philadelphia, PA.
- Baayen, R. H. and Smolka, E. (2020). Modelling morphological priming in German with naive discriminative learning. *Frontiers in Communication, section Language Sciences*. preprint on PsyArXiv, doi:10.31234/osf.io/nj39v.
- Box, G. E. P. (1976). Science and statistics. *Journal of the American Statistical Association*, 71:791–799.
- Bozic, M., Marslen-Wilson, W. D., Stamatakis, E. A., Davis, M. H., and Tyler, L. K. (2007). Differentiating morphology, form, and meaning: Neural correlates of morphological complexity. *Journal of cognitive neuroscience*, 19(9):1464–1475.
- Breiman, L. et al. (2001). Statistical modeling: The two cultures (with comments and a rejoinder by the author). *Statistical Science*, 16(3):199–231.
- Chersi, F., Ferro, M., Pezzulo, G., and Pirrelli, V. (2014). Topological self-organization and prediction learning support both action and lexical chains in the brain. *Topics in cognitive science*, 6(3):476–491.
- Choi, H.-B. (1929). Uli Mal Bon (Korean Grammar). *Seoul: Chom-Um-5a, 6th edition*.
- Chuang, Y. Y. and Baayen, R. H. (2021). Discriminative learning and the lexicon: NDL and LDL. In *Oxford Research Encyclopedia of Linguistics*. Oxford University Press (to appear).
- Chuang, Y.-Y., Bell, M. J., Banke, I., and Baayen, R. H. (2020a). Bilingual and multilingual mental lexicon: a modeling study with linear discriminative learning. *Language Learning*, pages 1–73.
- Chuang, Y.-Y., Loo, K., Blevins, J. P., and Baayen, R. H. (2020b). Estonian case inflection made simple. A case study in Word and Paradigm morphology with Linear Discriminative Learning. In Körtvélyessy, L. and Štekauer, P., editors, *Advances in Morphology*, pages 119–141. Cambridge University Press.
- Creemers, A. (2019). Data for: Opacity, transparency, and morphological priming: A study of prefixed verbs in Dutch.
- Creemers, A., Davies, A. G., Wilder, R. J., Tamminga, M., and Embick, D. (2020). Opacity, transparency, and morphological priming: A study of prefixed verbs in dutch. *Journal of Memory and Language*, 110:104055.
- Ferro, M., Marzi, C., and Pirrelli, V. (2011). A self-organizing model of word storage and processing: implications for morphology learning. *Lingue e linguaggio*, 10(2):209–226.

- Fon, J. (2004). A preliminary construction of Taiwan Southern Min spontaneous speech corpus. Technical Report NSC-92-2411-H-003-050-, National Science Council, Taipei, Taiwan.
- 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.
- Heitmeier, M. and Baayen, R. H. (2020). Simulating phonological and semantic impairment of English tense inflection with Linear Discriminative Learning. *The Mental Lexicon*, 15(3):384–421.
- Kuperman, V., Schreuder, R., Bertram, R., and Baayen, R. H. (2009). Reading of multimorphemic Dutch compounds: Towards a multiple route model of lexical processing. *Journal of Experimental Psychology: HPP*, 35:876–895.
- Landauer, T. and Dumais, S. (1997). A solution to Plato’s problem: The latent semantic analysis theory of acquisition, induction and representation of knowledge. *Psychological Review*, 104(2):211–240.
- Lee, H. B. (1989). *Korean grammar*. Oxford University Press, New York.
- Lewis, G., Solomyak, O., and Marantz, A. (2011). The neural basis of obligatory decomposition of suffixed words. *Brain and language*, 118(3):118.
- Luo, X., Chuang, Y.-Y., and Baayen, R. H. (2021). Judiling: an implementation in Julia of Linear Discriminative Learning algorithms for language modeling. <https://megamindhenry.github.io/JudiLing.jl/stable/>.
- Marantz, A. (2013). No escape from morphemes in morphological processing. *Language and Cognitive Processes*, 28(7):905–916.
- Martin, S. E. (1960). *Korean reference grammar*. American Council of Learned Societies.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., and Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.
- Milin, P., Feldman, L. B., Ramscar, M., Hendrix, P., and Baayen, R. H. (2017). Discrimination in lexical decision. *PLOS-one*, 12(2):e0171935.
- Mitchell, J. and Lapata, M. (2008). Vector-based models of semantic composition. In *ACL*, pages 236–244.
- Plag, I., Homann, J., and Kunter, G. (2017). Homophony and morphology: The acoustics of word-final S in English. *Journal of Linguistics*, 53(1):181–216.
- Shafaei-Bajestan, E., Tari, M. M., and Baayen, R. H. (2021). LDL-AURIS: Error-driven learning in modeling spoken word recognition. *Language, Cognition and Neuroscience*.

- Smolka, E. and Eulitz, C. (2018). Psycholinguistic measures for german verb pairs: Semantic transparency, semantic relatedness, verb family size, and age of reading acquisition. *Behavior Research Methods*, 50(4):1540–1562.
- Smolka, E., Zwitserlood, P., and Rösler, F. (2007). Stem access in regular and irregular inflection: Evidence from German participles. *Journal of Memory and Language*, 57(3):325–347.
- Sohn, H. M. (2013). *Korean*, volume 4. Korea University Press.
- Stockall, L. and Marantz, A. (2006). A single route, full decomposition model of morphological complexity: MEG evidence. *The Mental Lexicon*, 1(1):85–123.
- Taft, M. (2004). Morphological decomposition and the reverse base frequency effect. *The Quarterly Journal of Experimental Psychology*, 57A:745–765.
- Tomaschek, F., Plag, I., Ernestus, M., and Baayen, R. H. (2019). Modeling the duration of word-final s in english with naive discriminative learning. *Journal of Linguistics*, 57(1):123–161.
- Wood, S. N. (2017). *Generalized Additive Models*. Chapman & Hall/CRC, New York.
- Yang, B. S. (1994). *Morphosyntactic phenomena of Korean in role and reference grammar: psych-verb constructions, inflectional verb morphemes, complex sentences, and relative clauses*. PhD thesis, State University of New York, Buffalo.