

Making sense of spoken plurals

Elnaz Shafaei-Bajestan¹, Peter Uhrig² and R. Harald Baayen¹

¹Eberhard Karls Universität Tübingen | ²Friedrich-Alexander-Universität Erlangen-Nürnberg

Distributional semantics offers new ways to study the semantics of morphology. This study focuses on the semantics of noun singulars and their plural inflectional variants in English. Our goal is to compare two models for the conceptualization of plurality. One model (FRACSS) proposes that all singular-plural pairs should be taken into account when predicting plural semantics from singular semantics. The other model (CCA) argues that conceptualization for plurality depends primarily on the semantic class of the base word. We compare the two models on the basis of how well the speech signal of plural tokens in a large corpus of spoken American English aligns with the semantic vectors predicted by the two models. Two measures are employed: the performance of a form-to-meaning mapping and the correlations between form distances and meaning distances. Results converge on a superior alignment for CCA. Our results suggest that usage-based approaches to pluralization in which a given word's own semantic neighborhood is given priority outperform theories according to which pluralization is conceptualized as a process building on high-level abstraction. We see that what has often been conceived of as a highly abstract concept, [+PLURAL], is better captured via a family of mid-level partial generalizations.

Keywords: plural semantics, distributional semantics, FRACSS, Cosine Class Average, sound-meaning mapping, spoken word recognition

1. Introduction

Distributional semantics (Boleda, 2020; Firth, 1968; Günther, Rinaldi, & Marelli, 2019; Harris, 1954; Landauer & Dumais, 1997; Mikolov, Chen, Corrado, & Dean, 2013; Wang, Wang, Chen, Wang, & Kuo, 2019) offers new opportunities for understanding the semantics of word formation (see, e.g. Kisselew, Padó, Palmer, & Šnajder, 2015; Marelli & Baroni, 2015) and inflection (see, e.g., Baayen & Moscoso del Prado Martín, 2005). Given corpus-based semantic vectors (known

as embeddings in computational linguistics and natural language processing) for pairs of base words and corresponding complex words, several methods have been proposed that take as input the semantic vector of the base word, and that produce as output the vector of the complex word.

One such method is illustrated in Figure 1. Given the semantic vectors for two pairs of singulars and plurals (*table/tables* and *pen/pens*), and given the semantic vector for *banana* but no semantic vector for its plural, the semantic vector for *bananas* is obtained by first calculating the vectors that start at a singular and point to the corresponding plural (represented by blue vectors), and averaging these, resulting in an average shift vector (in red). This shift vector can then be applied to the vector of *banana*, resulting in the semantic vector for *bananas* (lower panel). Kisselew et al. (2015) calculated the average shift vector for each of a large set of German derivational affixes, and showed that this results in high-quality estimates of the meanings of derived complex words.

Another method for modeling word formation processes is offered by Marelli and Baroni (2015). They used matrix multiplication to obtain predicted semantic vectors for derived words, and showed that this method generated quantitative predictors that help explain variance in measures of lexical processing such as reaction times in visual lexical decision.

The abovementioned studies associate each derivational exponent with one conceptualization function that predicts the complex word from its base word. However, recent research on plural inflection has shown that the semantics of plurality can be too complex semantically to be adequately represented by a single shift vector. Chuang, Brown, Baayen, and Evans (2022) observed that for Russian, a fusional language, the shift vector for a singular to a plural noun embedding varies systematically by each of the six cases. Nikolaev, Chuang, and Baayen (2022) report that also for Finnish, an agglutinative language, plural shift vectors are different for each of the 14 cases. Furthermore, Shafaei-Bajestan, Moradipour-Tari, Uhrig, and Baayen (2022) report that in English, shift vectors are conditioned on the semantic class of the base word.

Figure 2 illustrates this phenomenon, highlighting three clusters of shift vectors in a two-dimensional plane constructed with the t-SNE dimension reduction technique (van der Maaten & Hinton, 2008). T-SNE is designed such that if observations cluster in a high-dimensional space, these clusters are (with high probability) optimally visible and distinct in its two-dimensional representation of this space. In Figure 2, words denoting people (e.g., *president*) are found predominantly in the lower right (in blue), words for communication (e.g., *language*) cluster in the center-left (in green), and words for animals (e.g., *rabbit*) are found in the upper right of the t-SNE plane (in red). The grey dots represent nouns that belong to other semantic classes. Shafaei-Bajestan et al.

(2022) broke down the set of nouns into some 400 semantic classes,¹ derived from WordNet (Miller, 1995). In this way, overlap between clusters is reduced, and plural vectors constructed by adding a class-specific shift vector to the vector of the singular are more similar to the empirical vectors. Following their approach, class-specific average vectors are visualized with arrows in Figure 2 for the three example clusters.

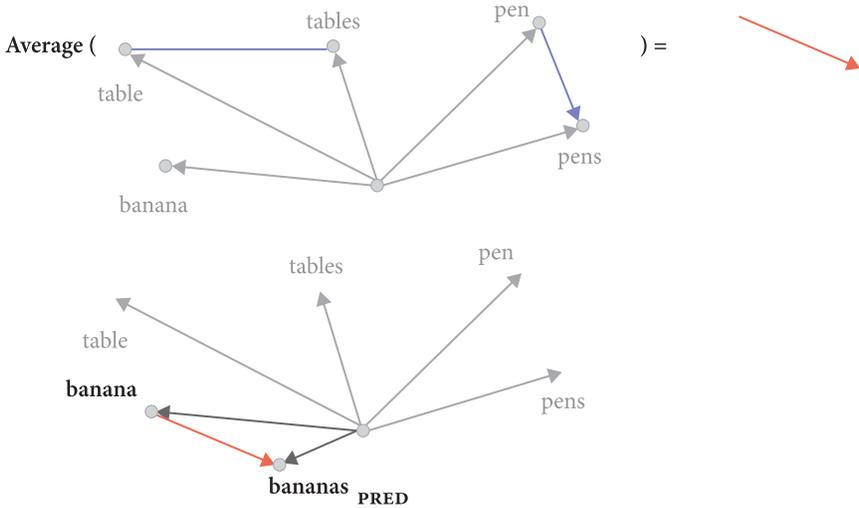


Figure 1. The average of the shift vectors for given singular-plural pairs (*table/tables*, *pen/pens*) is used to calculate the semantic vector of the unknown plural vector of *banana*

We note here that in a t-SNE plane constructed for singular and plural embeddings, plurals tend to occur very close to their singulars, and semantically similar words tend to be relatively close together. What the t-SNE analysis of the shift vectors shows is that how plural vectors are positioned with respect to their singular vectors varies considerably and systematically by semantic class.

In English, words belonging to a semantic class are not overtly marked by a specific exponent. However, there are languages that have grammaticalized class-conditional plural semantics. Swahili, a Bantu language, makes use of a large number of noun classes, many of which are semantically motivated (Polomé, 1967). For instance, the *m-wa* class is used for persons, the *ma-ji* class comprises words for objects that occur in clusters such as fruits, and the *ki-vi* class covers artifacts and tools. The Kiowa language from the Tanoan language family, has

1. See the complete list of classes at <https://osf.io/h6ck3/files/osfstorage/645fd97560fc5d2261ca098f>.

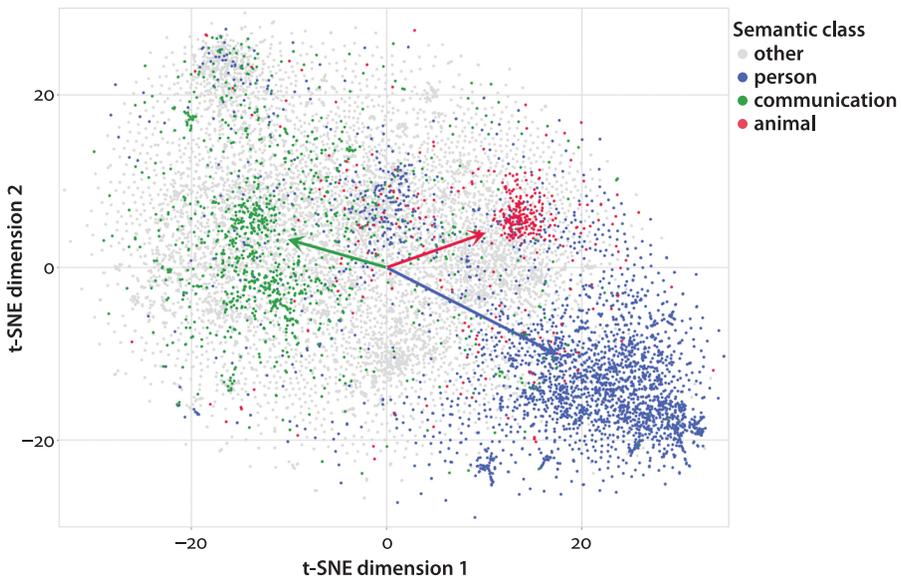


Figure 2. A projection of shift vectors (calculated from word2vec embeddings, Mikolov et al., 2013) onto a two-dimensional plane using the t-SNE dimension reduction algorithm reveals semantic clustering ($N=11749$). Each dot represents a shift vector, one for each singular-plural pair. Three semantic classes (based on WordNet) are highlighted: Animals (red), persons (blue), and communication words (green). The arrows represent the average shift vector for each of the three highlighted groups. Because there are many outliers far away from the main clusters, the arrows point towards the clusters but do not reach their centers. If the semantics of pluralization were identical for all words, a single cluster with its centroid at a distance from the origin would have emerged from the analysis. An interactive version of this plot can be found at <https://quantling.org/MentalLexicon/English/English-shift-space-w2v-3categories.html>

nine semantically motivated noun classes (Harbour, 2008, 2011). In contrast to English, Swahili and Kiowa have grammaticalized interactions of semantic class and plurality, indicating that speakers of these languages have made explicit in their grammars some perceived differences in what it means to be a plural.

Further evidence for the grammaticalization of number by semantic class is provided by languages with classifiers such as Mandarin Chinese (see, e.g. Yip & Rimmington, 2006). Constructions with numerals such as ‘two snakes’ require a classifier before the noun, as in *èr tiáo shé*, ‘two classifier snake’, and *èr zhī niǎo*, ‘two birds’. The classifier *tiao* is often found preceding long and flexible objects, such as snakes and rivers. By contrast, *zhī* is described as the classifier “for birds

and certain animals, one of a pair, some utensils, vessels, etc.”² Mandarin classifiers thus group multiples of objects together by their shape, size, or classes of artifacts (such as vehicles), among many others.

Grammaticalization of number by semantic class is perhaps unsurprising as for many concrete plural nouns, referents tend to occur in markedly different configurations. This is illustrated in Figure 3. Oranges and apples appear in groupings that are entirely different from the configurations in which cats and dogs are typically found. It is remarkable that configurational properties of multiples of objects are so well captured by word embeddings.



Figure 3. Multiples of different kinds of objects typically occur in different configurations. The images for oranges and dogs are obtained without modification from Vyagov (2021) and Gallice (2012). The images for apples and cats, obtained from Khursheed (2014) and Park (2013), have been cropped

The Kiowa noun system led Harbour (2008, 2011) to argue for a morphosemantic theory of number. The studies by Nikolaev et al. (2022); Chuang et al. (2022), and Shafaei-Bajestan et al. (2022) provide first steps in the formulation of such a theory. Focusing on noun inflection in English, the finding that plural semantics vary systematically by semantic class raises the question of how to formalize the conceptualization process that underlies the use of singular and

2. <https://chinese.yabla.com/chinese-english-pinyin-dictionary.php?define=zhi>.

plural nouns in English. Given a semantic vector v_{SG} for the singular of a given lexeme, we can ask what kind of conceptualization operation Φ is needed to map this singular onto the corresponding plural v_{PL} :

$$\Phi(v_{SG}) = v_{PL} \quad (1)$$

Shafaei-Bajestan et al. (2022) investigated two alternative implementations of Φ . The first implementation straightforwardly builds on the observed clustering in the original high-dimensional shift vector space.

$$\Phi(v_{SG}) = v_{SG} + v_{\text{shift|semantic class}} \quad (2)$$

For plurals that are ‘known’ to the algorithm, i.e. for which semantic vectors are available, a plural vector is semantically decomposed in the following way:

$$v_{PL} = v_{SG} + v_{\text{shift|semantic class}} + \epsilon \quad (3)$$

In other words, known plurals are composed of the singular semantic vector, a shift vector appropriate to the semantic class of their base words, and a residual error vector that specifies the semantics that are specific to a given plural form. As first pointed out by Sinclair (1991), regular inflected words often have their own collocational profiles, which are part of native speakers’ knowledge of the lexis of their language (see also Milin, Filipović Durdević, & Moscoso del Prado Martín, 2009; Moscoso del Prado Martín, Kostić, & Baayen, 2004). In (3), this word specific knowledge that is not predictable from the semantic class is represented by the error vector ϵ .

For ‘out-of-vocabulary’ plurals, i.e., plurals on which the conceptualization function Φ has not been trained, the semantic vector predicted for these plurals simplifies to:

$$v_{PL} = v_{SG} + v_{\text{shift|semantic class}} \quad (4)$$

This approach, which applies the approach to derivation of Kisselew et al. (2015) and extends it to the English plural, was laid out in Shafaei-Bajestan et al. (2022). Following their terminology, we henceforth refer to this method as Cosine Class Average (CCA).

An alternative approach to plural conceptualization builds on the FRACSS model of Marelli and Baroni (2015), which was originally developed for derivational morphology. When all singular semantic vectors are brought together as the row vectors of a matrix S , and all corresponding plural vectors are bundled together as the row vectors of a matrix P , a linear mapping M can be obtained by solving

$$SM = P.$$

Given M , the conceptualization function Φ is given by

$$\Phi(v_{SG}) = v_{SG} M, \quad (5)$$

the embedding of a plural is decomposed as

$$v_{PL} = v_{SG} M + \epsilon, \quad (6)$$

and for plurals for which an embedding is not available, the plural vector is estimated by

$$v_{PL} = v_{SG} M \quad (7)$$

The conceptualization functions Φ_{CCA} and Φ_{FRACSS} implement two different ways of thinking about the conceptualization of plural nouns. The perspective taken by Φ_{FRACSS} is global in nature: to predict a plural meaning, one needs to take into account the systematicities that govern all singular and plural pairs. In contrast, the perspective taken by Φ_{CCA} is local in nature, the hypothesis being that it is only the systematicities between singulars and plurals within the semantic class to which the lexeme belongs that are relevant.

However, although Φ_{CCA} straightforwardly implements the structure that is revealed by the t-SNE analysis to be present in the high-dimensional space of embeddings, precisely because it builds on semantic classes, it is actually a complex theory: semantic classes are not straightforward to define, and we therefore built on previous work done within the WordNet project (Miller, 1995). FRACSS, by contrast, is a theory that takes as its point of departure that there is a single plural operation. Because linear transformations are quite powerful, it is possible that the linear transformation of FRACSS will actually capture in its stride the semantic conditioning that is revealed by the t-SNE analysis.

Shafaei-Bajestan et al. (2022) observed that the plural vectors predicted by Φ_{FRACSS} and Φ_{CCA} are in general very similar. The former plural vectors were somewhat more similar to the word2vec gold standard plural vectors in terms of angle, the latter vectors were somewhat more similar in terms of vector length. In order to assess the relative merits of the two conceptualization functions, Shafaei-Bajestan et al. therefore investigated the relation between plurals' semantics and their forms. Different approaches to exploring this relation have been proposed in the literature. For example, Shillcock, Kirby, McDonald, and Brew (2001); Tamariz (2008), and Monaghan, Shillcock, Christiansen, and Kirby (2014) compared distances in the form space to distances in a distributional semantic space. Levy et al. (2021) employed methods from network science to build a multi-layer network that connects a phonological network to a semantic network. Lexical measures such as orthography-semantics consistency (Marelli, Amenta,

& Crepaldi, 2015; Siegelman et al., 2022) and phonologysemantics consistency (Amenta, Marelli, & Sulpizio, 2017) are now available for capturing the effects of form-meaning associations on lexical processing. Shafaei-Bajestan et al. used a different method, described in the following paragraph.

Shafaei-Bajestan et al. (2022) addressed the systematicities between form and meaning of English plural nouns using the Discriminative Lexicon Model (DLM) proposed by Baayen, Chuang, Shafaei-Bajestan, and Blevins (2019). This computational model approximates comprehension with a linear mapping from numeric form representations to semantic vectors. A core assumption of the DLM is that the form space and the semantic space are so similar and well calibrated to each other that simple linear mappings are sufficient for obtaining high-quality (albeit not perfect) predictions for meanings given forms (in comprehension), and for forms given meanings (in production). Given that according to this model the mappings between form and meaning are extremely simple, the quality of the model's predictions is critically dependent on the representations of form and of meaning. Within this modeling framework, it is therefore important to determine which of the available methods produces semantic vectors that are best aligned with the form vectors. Using form representations that encode which triphones are present in a word's form, Shafaei-Bajestan et al. found that plural vectors obtained with Φ_{CCA} could be predicted more precisely from their forms than plural vectors obtained with Φ_{FRACSS} , especially when the model was tested on held-out data. However, form vectors derived exclusively from (tri)phones do not do justice to either the richness or the enormous variability of spoken words.

The goal of the present study is not to provide improved semantic representations for plurals, but rather to address the shortcoming of Shafaei-Bajestan et al. (2022), where phone-based form representations were used. In the present study, we will replace phones with less abstract measurements to explore how well Φ_{CCA} and Φ_{FRACSS} are aligned with the audio signal itself. In what follows, we report two experiments addressing this goal. The first experiment follows Shafaei-Bajestan, Moradipour-Tari, Uhrig, and Baayen (2021) and makes use of a version of the Discriminative Lexicon Model that was developed specifically for auditory comprehension: the LDL-AURIS model. As this model derives form vectors for word tokens from the audio files of these tokens, it is well-suited to addressing the relation between words' speech audio and their meanings. The second experiment follows in the footsteps of Shillcock et al. (2001); Tamariz (2008), and Monaghan et al. (2014). These studies make use of distance measures, which for words' forms are based on orthographic or phonebased representations. In contrast, the present study uses distances between form vectors extracted from the audio files of word tokens, and compares these with the distances between semantic vectors for the corresponding word types. Our question of interest is

whether distances in ‘audio space’ are more similar to distances in semantic space when plural meanings are estimated with Φ_{FRACSS} or with Φ_{CCA} .

In what follows, we first introduce the dataset on which our analysis is based. We then report the study using LDL-AURIS (Section 3), followed by the study using distances (Section 4). We conclude with a discussion in Section 5.

2. Data

Our data are taken from the NewsScape English Corpus 2016 (Uhrig, 2018, 2021). This corpus is based on the US-American English-language recordings in the UCLA Library Broadcast NewsScape. The corpus contains TV news, including traditional newscasts and political talk shows as well as late-night shows and daytime talk shows. The vast majority of the language recorded can thus be assumed to be Standard English with a Network English accent. Overall, male speakers contribute much more than half of the words. In a sample of shows from broadcaster CNN for which transcripts were available online, about 80% of the words were uttered by male speakers (Uhrig, 2021). In terms of age distribution, younger speakers, especially children and teenagers, are underrepresented. The initial audio quality is usually quite good because much of the audio comes from recording studios, but the corpus also includes footage from outside reporters. There are also programs with audiences, who cheer and applaud, all reducing the quality of the audio signal.

For the year 2016, these recordings amount to more than 35,000 hours with around 269 million tokens of subtitles. Directly after subtitle extraction, the recordings are compressed to 240 MB/hour, using H.264 for the video stream and the Fraunhofer FDK library to produce an AAC audio stream at 96 kbit/sec. The subtitles are processed with Stanford CoreNLP (Manning et al., 2014) version 3.7.0, i.e. annotated for part of speech with the Penn Treebank tagset using CoreNLP’s caseless model.

Subtitles are not intended to be accurate transcripts of the words spoken. For the sake of readability, they usually omit false starts, often leave out words and phrases – in particular in fast dialogue – and are not designed to cope with overlaps. Commercials often do not have subtitles at all. We thus expect a certain error in the forced alignment, which is designed to align a transcript with the corresponding audio recording. Due to the transcript quality and the length of the recordings, most forced alignment systems are not suitable for this use case.

The phonemic transcriptions that are provided in the NewsScape English Corpus 2016 were obtained with a slightly modified version of Gentle (Ochshorn & Hawkins, 2015), which markets itself as a “robust yet lenient forced aligner

built on Kaldi³. Gentle runs an automatic speech recognition process with a bigram language model created from the transcript in the background. It is based on Kaldi ASR (Povey et al., 2011), which uses a purpose-built version of the CMUDict machine-readable pronunciation dictionary (<https://github.com/cmuspinx/cmudict>) that comes without stress information. Note that our dataset inherits the pronunciation variants offered by CMUDict, i.e. the word *either* will be represented either as [aɪðər] or as [i:ðər], depending on the pronunciation of the speaker. However, not all pronunciation variants typically found in American English are listed in the version of CMUDict used here. Thus *don't* is transcribed as [daʊnt] or as [daʊn], but vowel reductions that are typical of fast speech in combinations such as *don't know*, often rendered orthographically as *dunno*, are not covered. Thus we have to be aware that while our dataset offers some variation in pronunciation, which makes it more ecologically valid than purely dictionary-based approaches without speech recognition, it will contain instances where the transcript does not correspond to the spoken form because the transcription of the spoken form was not available in CMUDict.

Although the forced alignment of the corpus is reported to have a success rate of somewhere between 90% and 95% on average, we were able to select for inclusion only those files where Gentle's self-reported success rate was at least 97% because we only needed a much smaller dataset than the entire corpus. From this well-aligned part of the corpus, we selected the first 500 hours of programs in order of broadcast time and extracted a total of 750,816 audio word tokens of 2062 orthographic word form types that are either singular or plural nouns. A word form type was selected with the constraints that its token frequency was at least 70 in the 500-hour corpus and that a word2vec semantic vector was available for it. Proper names, plurals endings with anything other than an -s, and named entities were excluded from the dataset.

3. Study 1: Assessing the isomorphism of form and meaning with LDL-AURIS

3.1 Representations

For our investigations into the relation between spoken noun singulars and plurals and their semantics, we set up three sets of 300-dimensional semantic vectors. The first set of vectors based on pre-trained word2vec (Mikolov et al., 2013) embeddings (for both singulars and plurals) was retrieved from Mikolov

3. <https://lowerquality.com/gentle/>

(2013, Jul 30).⁴ For the compilation of the other two sets, the vectors for singular words are identical to those of the first set. For set 2, the plural vectors were calculated using Φ_{CCA} , and for set 3, using Φ_{FRACSS} . The semantic classes for the Φ_{CCA} conceptualization function are taken from WordNet (Ciaramita & Johnson, 2003; Fellbaum, 1998; Miller, 1995). For further details on the granularity of the conceptual classes used, see Shafaei-Bajestan et al. (2022).

To represent the speech signal of singular and plural nouns, we made use of the Continuous Frequency Band Summary Features (C-FBSF) developed by Shafaei-Bajestan et al. (2021). Figure 4 illustrates, for one example audio token which can be accessed through the QR link in the upper right of the figure, how these features are calculated. The algorithm takes the digital signal of a word as input (top panel) and divides it into temporal chunks delimited by the maxima of the periodic parts of the signal. The MEL spectrogram (middle panel) for each chunk is divided into 21 frequency bands, mirroring the frequency bands of the cochlea. For each frequency band of each temporal chunk, an order-preserving random sample of length 20 is taken from the list of log mel energy values and included in the C-FBSF vector (lower panel). Within each temporal chunk, correlation coefficients of the log mel energy values at the current frequency band with those of the following bands are calculated and appended. These numeric vectors are concatenated first within chunks and then between chunks.

The dimensionality of the auditory space is determined by the number of chunks in the longest word in the dataset, which is 13 for the present data. The algorithm stores 651 numbers per chunk. As a consequence, the auditory space is 8,463-dimensional. All C-FBSF vectors are zero-padded to match the 8,463 dimensions. For the audio token *apple* portrayed in Figure 4, which has a duration of 0.32 seconds, an initial 1953-dimensional vector is thus padded with 6510 trailing zeros.

3.2 Mapping audio vectors onto semantic vectors

The form vectors derived from the audio files were brought together as the row vectors of a form matrix C . The corresponding semantic vectors were bundled as the row vectors of three semantic matrices S : one with the word2vec gold standard vectors, one replacing word2vec plural vectors with FRACSS vectors, and one replacing word2vec plural vectors with CCA plural vectors.

4. In this study, we accept the embeddings of word2vec as ground truth. A question we leave to further research is whether multi-modal embeddings (see, e.g., Kiela, Bulat, & Clark, 2015; Kiela & Clark, 2017; Shahmohammadi, Lensch, & Baayen, 2021, for olfactory, auditory, and visual grounding of embeddings) will help enhance the modeling of the conceptualization of plurality in English.

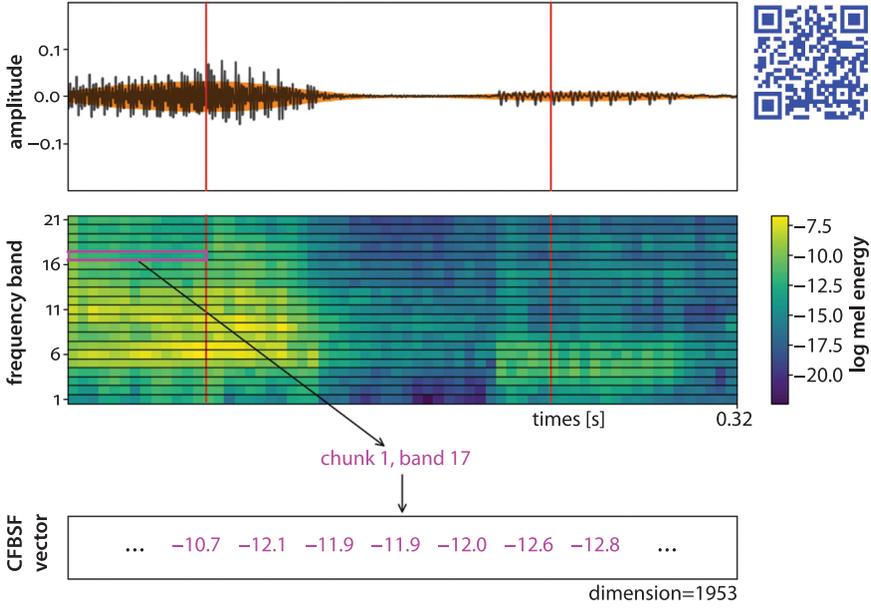


Figure 4. C-FBSF extraction for an audio rendition of the word *apple* (duration of 0.32 s). Click <http://go.redhenlab.org/pgu/0120> or scan the QR code to listen to the audio. The top panel presents the waveform in dark gray, and the smoothed Hilbert amplitude envelope of the signal in orange. The red vertical lines indicate chunk boundaries, located at the local maxima of the Hilbert envelope at 0.07 s and 0.23 s. The mid panel presents the corresponding log MEL spectrogram (using an FFT window size of 0.005 s) that depicts log energy values as colors over time (shown on the x-axis) split at 21 auditory filter banks according to the MEL scale (shown on the y-axis). The upper and mid panels are time-aligned. Lower panel: A segment of the output C-FBSF vector that corresponds to the first 7 sampled energy values from the 17th frequency band of the first chunk. The length of the C-FBSF vector compiled from all three chunks and 21 frequency bands before zero-padding is 1953

LDL-AURIS calculates a mapping F from the form matrix C to a semantic matrix S by solving

$$CF = S.$$

We used the normal equations (Faraway, 2005) for the multivariate multiple regression model to estimate F , using the Moore-Penrose algorithm (Moore, 1920) for matrix inversion. Once F has been estimated, we obtain the predicted semantic matrix \hat{S} as follows:

$$\hat{S} = CF \tag{8}$$

We write \hat{S} rather than S because, just as in regression one cannot draw a straight line through all the data points in an ellipsoid cluster, the predicted row vectors \hat{s} of \hat{S} cannot be identical to the gold standard row vectors s of S . A specific predicted semantic vector \hat{s}_i is judged to be accurate when it is closer to the corresponding gold-standard vector s_i than to any other semantic vector s_j , $j \neq i$. As a measure of semantic proximity, LDL-AURIS uses the Pearson correlation.

3.3 Training data and test data

To evaluate the quality of a mapping F , we split the data described in Section 2 into training and test sets using a stratified 10-fold cross-validation design with respect to the word form types. Cross-validation is a standard procedure in machine learning for model evaluation. In 10-fold cross-validation, the dataset is split into 10 parts (henceforth folds). Ten models are fitted on nine of the folds, each with a different fold held out, and are tested on the remaining fold. Stratification ensures that each fold is appropriately representative of the whole dataset. We made sure that all word form types appear in every training and test set and that the relative frequencies of audio tokens for each word form type are preserved across training and testing. Note that the plural form of a singular word or the singular form of a plural word might be absent from the data. Table 1 reports the number of word form types and the number of audio tokens averaged over the 10 folds for singular and plural words, broken down by whether the corresponding plural or the corresponding singular appears in the training set.

Table 1. The number of word form types and the average number of audio tokens (and standard deviation) across 10 stratified folds

Dataset	Word form types	Average audio tokens (SD)
Train	2062	675734 (1.26)
singular with plural	475	273893 (6.26)
singular without plural	979	243630 (9.10)
plural with singular	499	141214 (7.95)
plural without singular	109	16996 (4.81)
Test	2062	75082 (1.26)
singular with plural	475	30433 (6.26)
singular without plural	979	27070 (9.10)
plural with singular	499	15690 (7.95)
plural without singular	109	1888 (4.81)

3.4 Procedure

To evaluate the three kinds of predicted plural semantic vectors, we constructed the gold standard semantic matrices S_{CCA} , S_{FRACSS} , and $S_{word2vec}$, each of which contains row vectors for both singulars and plurals. The CCA and FRACSS vectors for plurals were estimated with Equations (4) and (7), respectively. Using the form matrix C , we calculated three mappings F_{CCA} , F_{FRACSS} , and $F_{word2vec}$. Applying (8), we then obtained the estimated vectors \hat{S}_{CCA} , \hat{S}_{FRACSS} , and $\hat{S}_{word2vec}$.

It is worth keeping in mind that the rows of the form matrix C are all distinct, and that the total number of rows in the C matrix equals the number of audio tokens in our dataset. In contrast, many rows of the semantic matrix S are duplicated, with as many duplicates for a given noun type as there are tokens for that noun type. As the estimated semantic vectors are *estimates*, the row vectors of \hat{S} are again all unique. However, if the mappings are of good quality, the estimated semantic vectors for the audio vectors of the same word type are expected to be highly similar.

We also note that because mappings are estimated for singulars and plurals jointly, changing the semantic vectors of plurals will affect the value of F in Equation (8) and will therefore affect the semantic vectors predicted for singulars. In other words, even though the form and meaning vectors of singulars are identical across the three datasets, the estimated singular vectors will be somewhat different, depending on the quantitative properties of the plural semantic vectors.

We evaluated the quality of the estimated semantic vectors in two ways. First, we calculated accuracy on both the training data and the test data. The top- N accuracy of the model is defined as the percentage of the test items for which the predicted vector is among the top- N vectors in terms of strength of correlation with the targeted output vector.

Since the word types have highly unequal frequencies with a Zipf-like probability distribution, we also used the F_1 score (the harmonic mean of precision and recall) to gauge the quality of the mappings F and the different semantic vectors for plurals that these mappings predict.

3.5 Results

3.5.1 Recognitions

Figure 5 presents top- N accuracy (for $N=1, 2, \dots, 5$) for training data (dark blue) and test data (light yellow), for Φ_{CCA} (left) and Φ_{FRACSS} (right). Unsurprisingly, accuracy increases as the criterion for accuracy is relaxed by increasing N . For $N=5$, Φ_{CCA} is successful for up to 25% of the training tokens and 22% of the test

tokens. For $N=1$, accuracy on training is 15% and around 13% for held-out data. The drop in accuracy for test data as compared to training data is as expected. What is important is that the reduction in accuracy is relatively small, an indication that the mappings are not severely overfitting the data.

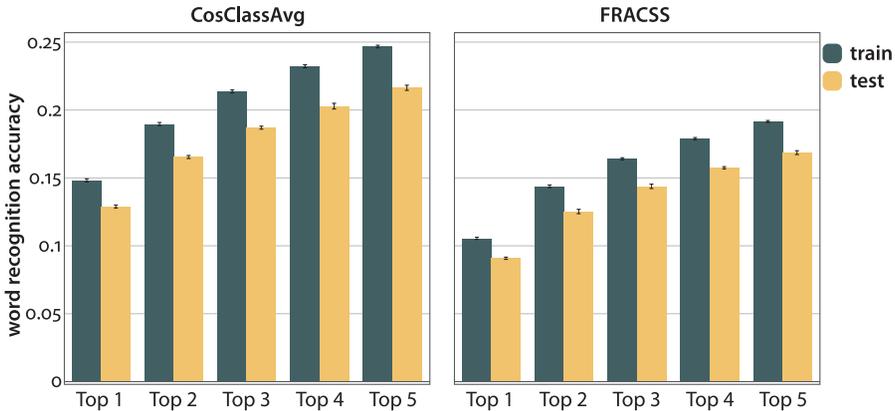


Figure 5. Top- N accuracy of word recognition on the training set (dark blue) and on the test set (light yellow) for CCA (left) and FRACSS (right). The small 1SE error bars indicate that there is very little variability across the 10 cross-validation folds

Recognition accuracies are low. This is unsurprising, for various reasons. First, it is worth noting that standard deep-learning speech recognition systems also do not perform well on the task of isolated word recognition.⁵ It is for the recognition of connected speech that deep learning models show truly impressive performance.

Second, the audio tokens used in the present study have all been cut from the corpus automatically based on the word boundaries identified during the forced alignment process (see Section 2). There are many instances in which these cut-points do not coincide with the exact word boundaries, even if the identification of such boundaries was possible at all.

Third, the models are trained and tested on auditory data spoken by numerous speakers, with no prior speaker normalization. The speakers come

5. Comparing the performance of LDL-AURIS with the performance of two Automatic Speech Recognition (ASR) systems, DeepSpeech2 and Kaldi, on 98,883 audio tokens from the NewsScape 2016 corpus, Shafaei-Bajestan et al. (2021) found that pre-trained Kaldi and pre-trained DeepSpeech2 performed at 22% and 20%, respectively, while LDL-AURIS trained on far fewer data performed at 12%.

from different age groups and genders in diverse environments, some of which have background noise.

For the purposes of the present study, the relative performance for plural semantic vectors of the Φ_{FRACSS} and Φ_{CCA} is at the center of interest. A comparison of the left and right panels of Figure 5 shows that Φ_{CCA} outperforms Φ_{FRACSS} for both training and test data.

Figure 6 spotlights the top 1 accuracy measure across 10 cross-validation folds evaluated on the training sets and on the test sets for the three mappings F_{CCA} , F_{FRACSS} , and F_{word2vec} , using box and whiskers plots. The mapping using CCA reaches a median accuracy of 14.91% on the training sets and 12.95% on the test sets, very close to the mapping that uses the semantic vectors provided by word2vec for plurals (15.05% and 12.95% for training and test respectively). The mapping using FRACSS plural vectors falls behind the CCA plural vectors by a 4.34 and a 3.83 percentage-point difference in median accuracy on the training and the test sets, respectively.

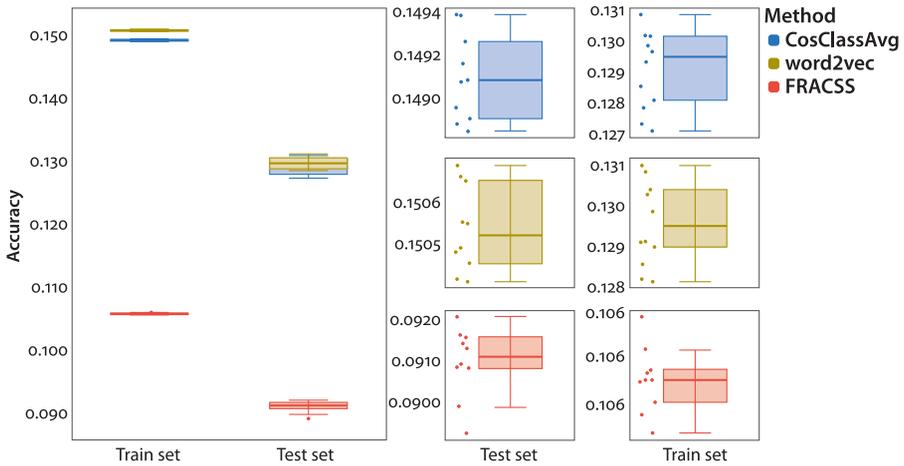


Figure 6. Comparison of the accuracy of CCA (blue), word2vec (light yellow), and FRACSS (red), evaluated with a stratified 10-fold cross-validation design on the training and the test sets. The left panel shows the box plots for all datasets crossed by methods. The panels on the right zoom in on individual plots for improved readability. Each data point represents one fold

As shown in Figure 7, when the performance of the mappings is evaluated with weighted average F_1 scores,⁶ the mapping using CCA outperforms the mapping using word2vec, which in turn outperforms mappings based on FRACSS on the training and the test sets of all folds.

We also compared the accuracy of these mappings with mappings that were set up with semantic matrices for which the row vectors were randomly permuted, breaking the relationship between words' forms and meanings. As variability in model performance turned out to be very small (see Figure 5), models were trained and tested on just the first fold. The reduction in accuracy and F_1 scores on the training and the test sets are reported in Table 2. Consistent with the previous results, the reduction in performance under random permutation is larger for CCA as compared to FRACSS for all dataset-metric pairs.

As a final check on the quality of the two implementations of plural conceptualization, we examined accuracy for four subsets of the data, defined by the combinations of the singular and plural forms that were available to the model for training. Figure 8 shows that accuracy was highest for singular tokens the plural token of which also was encountered during training. The group with the second-highest accuracy comprises plural tokens that have a singular token in the training set. Accuracies are substantially lower for singular and plural tokens that were not accompanied by a corresponding plural or singular token during training. (For the counts of auditory tokens within each group, see Table 1 above). For both conceptualization methods, recognition accuracy benefits from training on tokens of both numbers. Comparing CCA (left panel) with FRACSS (right panel), we can conclude that CCA accuracies are greater than those for FRACSS, except for held-out plurals that have no corresponding singulars in the training dataset.

Table 2. Reduction in accuracy and in weighted averaged F_1 score of LDL-AURIS on one cross-validation fold when the relation between form and meaning is broken by randomly permuting the rows of the semantic matrix

Model	Reduction in accuracy		Reduction in F_1 score	
	Training	Test	Training	Test
FRACSS	0.091	0.088	0.074	0.073
CCA	0.135	0.126	0.093	0.088

6. A weighted averaged F_1 score is the average of the F_1 scores for all word types weighted by the number of true instances for each word type.

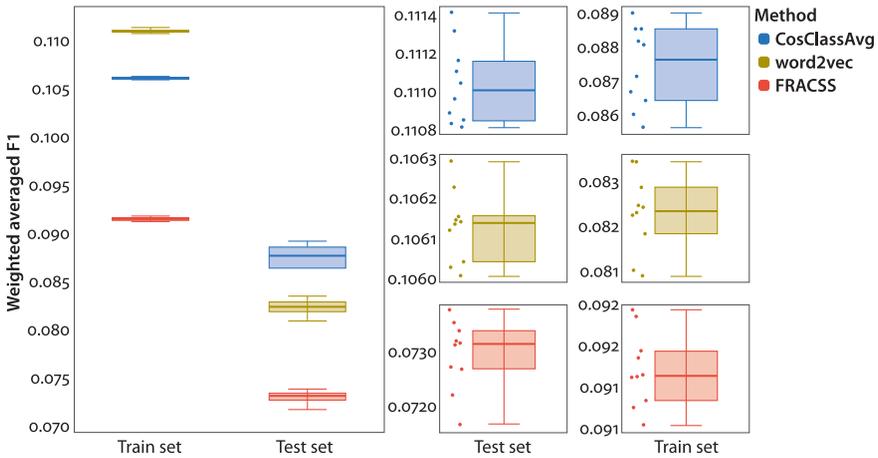


Figure 7. Comparison between the weighted average F_1 scores for the mappings using CCA (blue), word2vec (light yellow), and FRACSS (red), evaluated with a stratified 10-fold crossvalidation design on the training and the test sets

In summary, higher accuracies and higher higher average F_1 scores were obtained with CCA compared to FRACSS, across both training and test data.

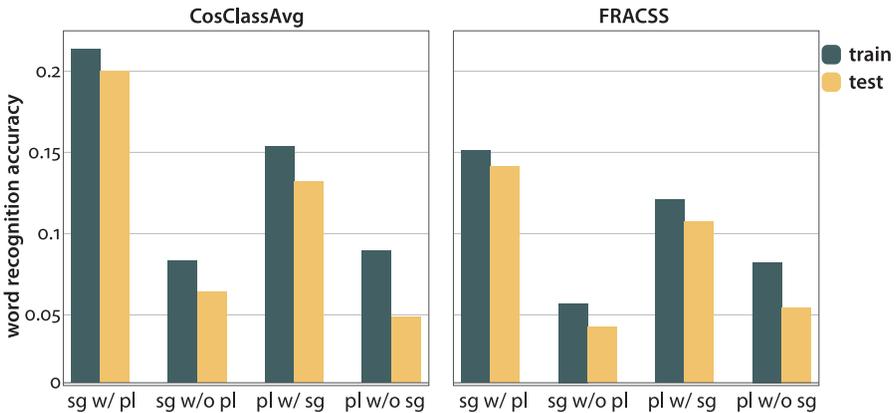


Figure 8. Within-group accuracy for different word groups. In group names, sg stands for *singular*, pl for *plural*, w for *with* and w/o for *without*. A word in the *sg w/o pl* group is a singular word whose plural form is not in the training set. CCA accuracies are greater than those for FRACSS, except for held-out plurals that have no corresponding singulars in the training dataset

3.5.2 Errors

The confusion matrices in Figure 9 summarize the predictions of the models for the items in the training set.⁷ Rows correspond to the number feature for the target word: singular words are counted in the first row and plural words in the second row of each matrix. The columns group words by the four combinations of number and whether the lexeme is properly predicted. Errors are highlighted in red and correct recognitions are highlighted in blue. In addition to the absolute number of tokens, the percentage with regard to the sum of the row is given in each cell. Training with plural vectors obtained with CCA gives rise to larger numbers of correctly predicted target words in the plural as well as in the singular group.

The two models rarely confuse number when they correctly predict the lexeme. However, when a lexeme is not recovered correctly, the two models make remarkably different kinds of errors. For plural vectors generated with CCA, the predicted and targeted word types tend to match in terms of number. In contrast, when plural vectors are generated using FRACSS, the model mostly predicts plural forms irrespective of the target word's number.

Overall, 81% of the predictions of the mapping using CCA match the target in terms of number. This percentage reduces to only 50% for models with plurals generated by FRACSS.⁸ Apparently, FRACSS introduces a bias for plural semantics. From a methodological perspective, these findings point to the importance of considering predicted plural vectors not just in isolation, but within the full noun system that in addition to plurals also contains singulars.

Apparently, plural vectors generated by Φ_{CCA} are better aligned with information about number in the audio signal than is the case for plural vectors generated by Φ_{FRACSS} . This suggests that the global perspective of FRACSS, which takes all singular-plural pairs into account, rather than just those singular-plural pairs that belong to a conceptual class and that share a shift vector, is suboptimal.

It is surprising that the FRACSS method, which is quite powerful as the underlying linear algebra allows for various geometric modifications such as scaling (stretching or shrinking) and rotation, performs somewhat less well than CCA, a method that from a geometric perspective only implements shifts between vectors. On the other hand, CCA needs to be informed about semantic classes and assumes that these are given, but setting up these classes is not trivial.

7. The results on the test set are very similar and reported in the supplementary materials.

8. The estimated difference in proportions between groups at 0.308 is supported by a proportions test (95% confidence interval of 0.306 and 0.309, $\chi^2(1) = 140897$, $p \ll .0001$).

		Predicted by CosClassAvg			
		SG same lexeme	PL same lexeme	SG other lexeme	PL other lexeme
Actual	SG	77735 (15.0%)	5184 (1.0%)	361848 (69.9%)	72762 (14.1%)
	PL	6185 (3.9%)	22860 (14.5%)	47348 (29.9%)	81800 (51.7%)

		Predicted by FRACSS			
		SG same lexeme	PL same lexeme	SG other lexeme	PL other lexeme
Actual	SG	53308 (10.3%)	17777 (3.4%)	132915 (25.7%)	313529 (60.6%)
	PL	946 (0.6%)	18018 (11.4%)	7158 (4.5%)	132082 (83.5%)

Figure 9. Confusion matrices for the performance of LDL-AURIS with CCA (top) and FRACSS (bottom) vectors on the training set. Correct predictions appear in blue and wrong predictions appear in red. Percentages are normalized to the sum of rows.

$N_{SG} = 517, 529$, $N_{PL} = 158, 204$ for both matrices

3.6 An alternative approach

The models so far were built on the assumption that error vectors ϵ in Equations (3) and (6) are all measurement noise and reducible. Thus, audio-to-semantic mappings were trained *and* evaluated using plurals predicted by the CCA or the FRACSS conceptualization methods, which build plural vectors without the error term.

Alternatively, one might assume that the error vectors in Equations (3) and (6) are entirely word-specific semantic information and irreducible.⁹ In this case, it would be appropriate to train the audio-to-semantic mappings using the plurals predicted by CCA or FRACSS *but* to evaluate the mappings using the corpus-extracted word2vec vectors, which include the error term. Accordingly, we evaluated the mappings by employing the corpus-extracted word2vec embeddings as gold standard, and comparing the gold standard with semantic vectors predicted

9. We actually believe that the error term comprises both measurement noise and word-specific knowledge.

by the previously obtained form-to-meaning mappings F_{CCA} and F_{FRACSS} . For this evaluation, we used the data of Section 3.3. The word2vec vectors are brought together as the rows of a matrix $S_{word2vec}$. We estimated plural vectors for the test set for each of the ten models described in Section 3.3, and defined a plural vector to be accurate if it was most correlated with the corresponding vector of $S_{word2vec}$.

Accuracies obtained from the second evaluation approach are reported in Figure 10. They are slightly higher for FRACSS for Top 1 evaluation, and slightly higher for CCA for Top 2–5 evaluation. Both of the trained mappings get approximately equally close to the empirical word2vec plural vectors. This suggests that neither CCA nor FRACSS plural vectors are very different from word2vec plural vectors. The results provide further evidence for the observation by Shafaei-Bajestan et al. (2022) that CCA and FRACSS vectors are generally similar to word2vec vectors. In their study, plural vectors predicted by CCA and FRACSS were compared with corpus-extracted plurals using cosine similarity as the measure (a median cosine similarity of 0.71 for CCA and 0.75 for FRACSS).

These accuracies are within the range of accuracies observed in Figure 5, which were obtained with CCA and FRACSS vectors as gold standard for evaluation. From this analysis, we conclude that the CCA, FRACSS, and empirical vectors are all of very similar, and probably high quality. The low word recognition accuracy is not due to a lack of quality of these vectors, but to the enormous variability in the speech signal and the noise in our data, which, as mentioned above, also challenges state-of-the-art deep learning systems.

4. Study 2: Assessing the isomorphy of form and meaning with distance measures

We have seen that within the framework of the DLM, by most of the measures discussed in Section 3, semantic vectors conceptualized with Φ_{CCA} enable a better mapping between words' audio and their meanings compared to Φ_{FRACSS} . In this section, we compare the two conceptualization functions with a method that does not depend on the assumptions of the DLM, including the assumption that the form space and the semantic space can be transformed into each other simply using linear mappings.

In what follows, we follow in the footsteps of Monaghan et al. (2014), and investigate possible parallelism between the form space and the meaning space using distance measures. Monaghan et al. reported that phone-based form similarity (gauged with Euclidean distance and edit-distance measures) correlated weakly ($r \approx 0.03$) yet significantly with semantic similarity (gauged with the

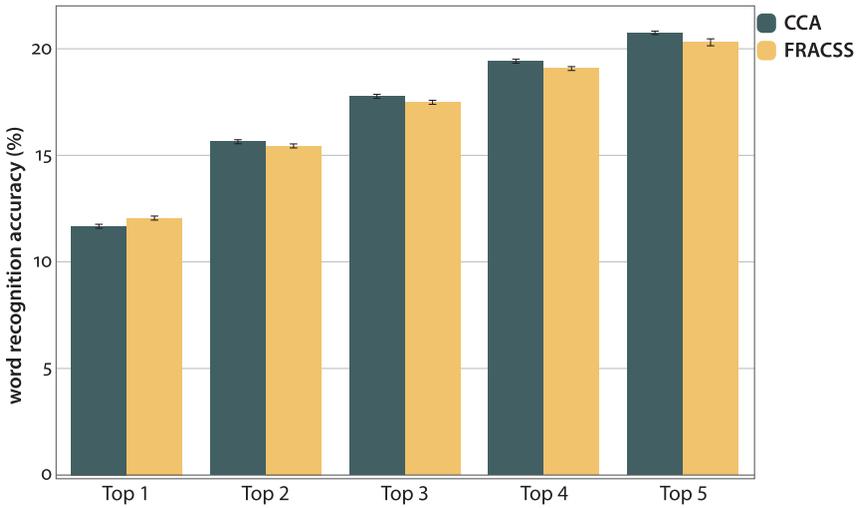


Figure 10. Top- N accuracy (%) of mappings from audio to semantics, trained either with CCA or FRACSS vectors, but evaluated with wordzvec vectors as gold standard. The small 1SE error bars indicate that there is very little variability across the 10 cross-validation folds

cosine similarity measure between contextual co-occurrence vectors). They observed correlations for monosyllabic words, uninflected lemmas, and monomorphemic polysyllabic words.

We calculated, for all pairs of words, both phone-based and audio-based measures of their form similarity. In parallel, we calculated the semantic similarity for these pairs of words, using both the original wordzvec vectors for plural nouns, as well as plural vectors generated with FRACSS and generated with CCA. For the calculations, the words from the test set of the first fold of the data (see Section 3.3) for which phonemic information was also available were used ($N=60, 206, V=2057$). Phone-based form similarities and semantic similarities were calculated for pairs of word form types while audio-based form similarities were calculated for pairs of audio tokens.

First, we calculated the Damerau-Levenshtein phone-based edit distance for all pairs of word form types in our dataset, which unlike the dataset of Monaghan et al. (2014), is restricted to singular and plural nouns. Thus, $\binom{2057}{2} = 2,114,596$

phone-based edit distance values were obtained. Second, the cosine distances¹⁰

10. The cosine distance zooms in on the angle between two vectors, ignoring differences in length. The greater the angle between two vectors, the greater their cosine distance is.

between the semantic vectors of these pairs of word form types were computed using three different semantic vectors for plurals: the word2vec plural vectors, the plural vectors obtained with FRACSS, and the plural vectors obtained with CCA. We then calculated the Pearson correlations between these phone-based edit distances on the one hand, and the cosine distances on the other hand.

Table 3. Correlations between phone-based and audio-based form distances and semantic distances for three different semantic spaces: CCA, word2vec, and FRACSS. For the audio-based analysis correlations averaged over many samples of audio tokens are reported, with SD and the number of trials N in which the correlation was significant at $\alpha = 0.05$

Form distances	Statistic	CCA	Word2Vec	FRACSS
phone-based	r	0.061	0.041	0.038
audio-based	\bar{r}	0.019	0.015	0.007
	SD	0.003	0.003	0.005
	N	1000	1000	908

The upper part of Table 3 presents the correlation r between the form and meaning distances for all pairs of word form types. These correlations are smallest for FRACSS, intermediate for word2vec, and greatest for CCA (all $p < .0001$). The correlations observed for vectors are of the same order of magnitude as the correlations reported by Monaghan et al. (2014). For a random permutation of the word2vec vectors, the correlation was reduced by a factor of 10 to 0.004. Our replication study shows, first, that there is indeed some isomorphism between the form space and the semantic space; and second, that the observations of Monaghan et al., which were based on uninflected words, generalize to plural inflection. Importantly, this analysis shows that, consistent with the results reported in the previous section, CCA-based semantic vectors provide stronger correlations between form distances and semantic distances than FRACSS-based semantic vectors.

To assess the isomorphy between the form space and the semantic space using the speech signal instead of phone-based representations, we again compared all pairs of words, assessing their semantic distance with the cosine distance, and evaluated the dissimilarity of their speech signals by means of the Euclidean distance between their C-FBSF form vectors, which were calculated from the

speech signal as described above in Section 3.1.¹¹ As can be seen in Figure 11, audio-based distance increases with phone-based edit distance. As expected, the variability in the audio-based distance as a function of phone-based distance is huge.

To avoid comparing audio tokens of the same word form type, which is likely to inflate similarity compared to a phone-based analysis in which types do not have non-identical replicates, we randomly sampled for each word type one audio token and then calculated the form distance for each pair of the selected tokens. Subsequently, the correlation between the form distances and the semantic distances of all 2, 114, 596 pairs of word types was computed. This procedure was repeated 1000 times. The lower part of Table 3 presents summary statistics for the audio-based comparisons of form and meaning. The average correlations across trials \bar{r} , the standard deviations SD, and the number of trials N , out of 1000 trials, for which the correlation was significant at $\alpha = 0.05$ are reported.

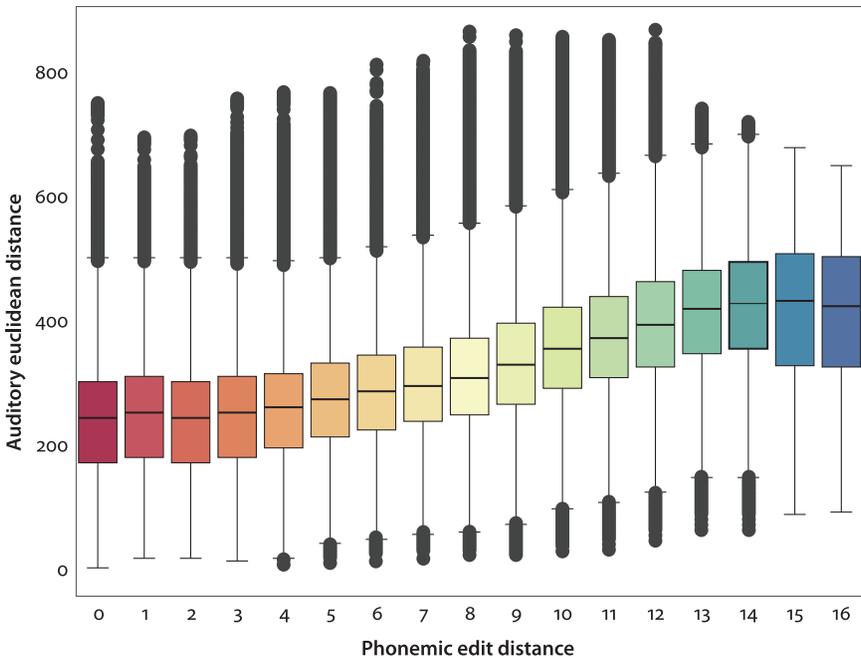


Figure 11. Auditory Euclidean distance as a function of phone-based edit distance

11. By way of example, the list of the top 10 nearest neighbors of an audio token of the word *positions* includes two other audio renditions of the same word form *positions*, one audio token of the singular form *position*, four tokens with similar word-initial phones (*protester*, *process*, *procedures*, and *perspective*), and three tokens with similar sounding word-final syllables (*information* and *conditions*).

Correlations are greatest when semantic vectors are estimated with CCA, and weakest when they are estimated with FRACSS. Randomization of the word2vec semantic vectors resulted in reduced correlations ($M=0.003$, $SD=0.003$). The number of trials in which a significant correlation was observed decreased from 1000 to 744.

The audio-based correlations are substantially lower than the phone-based correlations. This is unsurprising, as the phone-based analysis does not take into account the huge variability in the speech signal and the highly reduced forms that speakers use in daily speech (see, e.g., Johnson, 2004). The small magnitudes of these correlations show that evaluations of the similarities between form and meaning run the risk of overestimating these similarities if form representations are based on letters or phones. At the same time, it is surprising that significant correlations, however small, can be detected with this method for real, highly diverse, and noisy speech.

Importantly, our audio-based evaluation of the isomorphism between form and meaning provides further support for the possibility that plural semantic vectors conceptualized with Φ_{CCA} provide a better match with plural forms compared to Φ_{FRACSS} . Above, we showed that Φ_{CCA} also does better than Φ_{FRACSS} at predicting meaning from form, using a linear mapping within the framework of the DLM. The present observations show that this result is unlikely to be an artifact of this linear mapping. Plural vectors created with Φ_{CCA} better mirror similarities that exist in the form space.

5. General Discussion

Shafaei-Bajestan et al. (2022) showed that shift vectors, which represent the change in meaning when going from a noun's singular to its plural, cluster by the semantic class of the base word. Apparently, modeling the semantics of plurality in English with a single shift vector in semantic space, does not do justice to the intricacies of the English noun plural. Likewise, formal accounts of plurality based on abstract features such as [+PLURAL] lack precision. From a typological perspective, languages such as Swahili and Kiowa have grammaticalized semantic-class-based differences in their noun classes, and languages such as Mandarin Chinese, Japanese and Korean make use of classifiers that are often linked to semantic classes, whereas in English, systematic differences between words' referents are present only in the semantics as 'soft constraints' (see also Corbett, 2000).

The observations in the previous paragraph raise the question of how to account for the conceptualization of plurals. Do we have to construct different

plural shift vectors for each semantic class (CCA)? Or is it possible to predict a noun's plural meaning from a general rule that is informed by all singular and plural nouns in the language (FRACSS)?

Shafaei-Bajestan et al. (2022) addressed this question using the Discriminative Lexicon Model, and observed that semantic vectors for plurals generated with CCA were better aligned with triphone-based form vectors than plural vectors generated using FRACSS. The present study investigated whether this result generalizes to spoken language. In Section 3, using the form vectors calculated from the speech signal of a large number of audio tokens extracted from the NewsScape resource, we observed that the semantic vectors generated by CCA can be predicted with greater precision from their audio than is possible for plural vectors generated by FRACSS. This holds both for training data and test data. Furthermore, random pairing of form and meaning vectors degrades accuracy more for CCA than for FRACSS.

The superior alignment of the CCA vectors with the form space, compared to FRACSS, received further support in Section 4 from a series of analyses that do not depend on the linear mappings of the Discriminative Lexicon model, and instead follow up on a study by Monaghan et al. (2014). For each pair of word forms in our dataset, we calculated their distance from one another in semantic space as well as their distance from one another in form space (using both phone-based and audio-based distance measures). Correlations of the distances in form and meaning were stronger for plural vectors obtained with CCA compared to plural vectors obtained with FRACSS. We can conclude that approximating the conceptualization of noun plurals in English with CCA likely offers a more precise window on the processing of English noun singulars and plurals in the mental lexicon, the reason being that a better alignment of form and meaning implies greater regularity and hence facilitation of lexical processing.

It is surprising that FRACSS does not perform as well as CCA, as the linear mapping that FRACSS makes use of is more powerful than the simple vector addition that CCA employs. In fact, given the common assumption that pluralization is a unitary operation, a linear mapping as proposed by FRACSS makes sense: A global singular-meaning-to-plural-meaning mapping is informed by all the singular and plural nouns experienced by the language user, and can be expected to make optimal use of the shared semantic similarities that are supposed to exist between all singular nouns and between all plural nouns. However, careful inspection of the distributional semantics of English noun plurals shows that in this language, the semantics of pluralization is not as uniform as previously supposed. Conceptualization of English plurals is much more locally determined. It appears that, taking into account how nouns in

distant semantic classes are pluralized is therefore detrimental to prediction accuracy.

When comparing FRACSS with CCA, however, it should be noted that CCA takes for granted that information is available about which semantic class a lexeme belongs to. As semantic classes are not straightforward to induce in a bottom-up manner – we made use of WordNet – CCA is given much more refined information to work with compared to FRACSS. In terms of model complexity, CCA is therefore likely to be the more complex model. But this increased complexity appears to be justified by greater prediction accuracy.

We conclude by placing the present study in the wider context of research on spoken morphology. The data considered in our study are restricted to lexemes that take {S} as the exponent for plurality. Recent studies on English word-final /s/ have shown that the acoustic duration of /s/ varies systematically with the semantics realized with the /s/ (Plag, Homann, & Kunter, 2017; Tomaschek, Plag, Ernestus, & Baayen, 2019). In other words, an exponent that in standard theories is assumed to be realized in exactly the same way, independent of its function, is now known to be articulated in ways that reflect its function. One of the functions of /s/ considered in these investigations is the realization of plurality on nouns. The present study provides evidence that pluralization for English nouns is not a monolithic and general semantic operation, but varies with the semantic class of the base word. Importantly, we have shown that these differentiated plural semantics are aligned with words' audio signal. The way in which plurals are articulated therefore depends on the semantics of their base words. All that we have accomplished in the present study is demonstrate this fact. In what way these different class-specific semantics are realized in the speech signal is at present unknown, and is a topic for further research. A better understanding of the phonetic realization of English plurals may further our understanding of why deep learning models in natural language processing and artificial intelligence, which are superbly tuned to the distributional statistics of language use, are so successful. This, in turn, may contribute to enhancing models of lexical processing in the mental lexicon.

Funding

This research was made possible by funding from the ERC, project WIDE-742545, awarded to RHB, and by funding from the Competence Network for Scientific High Performance Computing in Bavaria to PU.

Acknowledgements

The authors are indebted to two anonymous reviewers, Jessica Nieder, and Melanie Bell for their critical and constructive feedback on our manuscript.

Authors' contributions

ESB and RHB conceptualized and designed the study. PU collected and processed the NewsScape data. ESB created the materials in consultation with PU. ESB implemented the simulations. All authors were involved in the data analyses and interpretation. All authors were involved in drafting the manuscript.

References

- Amenta, S., Marelli, M., Sulpizio, S. (2017). From sound to meaning: Phonology-to-Semantics mapping in visual word recognition. *Psychonomic Bulletin and Review*, 24 (3), 887–893. [doi](#)
- Baayen, R. H., Chuang, Y.-Y., Shafaei-Bajestan, E., 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*, 1–39. [doi](#)
- Baayen, R. H., & Moscoso del Prado Martín, F. (2005). Semantic density and past-tense formation in three Germanic languages. *Language*, 81, 666–698. [doi](#)
- Boleda, G. (2020). Distributional Semantics and Linguistic Theory. *Annual Review of Linguistics*, 6, 213–234. [arXiv:1905.01896v4](#). [doi](#)
- Chuang, Y.-Y., Brown, D., Baayen, R. H., Evans, R. (2022). Paradigm gaps are associated with weird “distributional semantics” properties: Russian defective nouns and their case and number paradigms. *submitted*. Retrieved from <https://psyarxiv.com/t7xba/download?format=pdf>. [doi](#)
- Ciaramita, M., & Johnson, M. (2003). Supersense tagging of unknown nouns in wordnet. *Proceedings of the 2003 conference on empirical methods in natural language processing* (p. 168–175). USA: Association for Computational Linguistics. [doi](#)
- Corbett, G. G. (2000). *Number* (S. R. Anderson et al., Eds.). Cambridge, UK: Cambridge University Press. [doi](#)
- Faraway, J. J. (2005). *Linear models with r*. Boca Raton, FL: Chapman & Hall/CRC. Retrieved from <http://www.stat.lsa.umich.edu/faraway/LMR/>
- Fellbaum, C. (1998). *WordNet: An electronic lexical database*. Cambridge, MA: MIT Press. [doi](#)
- Firth, J. R. (1968). *Selected papers of J. R. Firth, 1952–59*. Indiana University Press.
- Gallice, G. (2012). *Flickr – ggallice – street dogs (1)*. Wikimedia Commons. Retrieved 2022-5-31, from [https://commons.wikimedia.org/wiki/File:Flickr_-_ggallice_-_Street_dogs_\(1\).jpg](https://commons.wikimedia.org/wiki/File:Flickr_-_ggallice_-_Street_dogs_(1).jpg) (This file is licensed under the Creative Commons Attribution 2.0 Generic license.)
- Günther, F., Rinaldi, L., Marelli, M. (2019). Vector-Space Models of Semantic Representation From a Cognitive Perspective: A Discussion of Common Misconceptions. *Perspectives on Psychological Science*, 14 (6), 1006–1033. [doi](#)

- doi Harbour, D. (2008). *Morphosemantic Number: From Kiowa Noun Classes To UG Number Features* (1st ed.). Dordrecht: Springer.
- doi Harbour, D. (2011). Valence and atomic number. *Linguistic Inquiry*, 42 (4), 561–594.
- doi Harris, Z. S. (1954, 8). Distributional Structure. *WORD*, 10 (2–3).
- Johnson, K. (2004). Massive reduction in conversational American English. *Spontaneous speech: data and analysis. proceedings of the 1st session of the 10th international symposium* (pp. 29–54). Tokyo, Japan.
- Khursheed, O. (2014). *Apples of kashmir valley*. Wikimedia Commons. Retrieved 2022-5-31, from https://commons.wikimedia.org/wiki/File:Apples_of_Kashmir_valley.jpg (This file is licensed under the Creative Commons Attribution-Share Alike 4.0 International license.)
- doi Kiela, D., Bulat, L., Clark, S. (2015). Grounding semantics in olfactory perception. *ACL-IJCNLP 2015 – 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing of the Asian Federation of Natural Language Processing, Proceedings of the Conference*, 2, 231–236.
- doi Kiela, D., & Clark, S. (2017). Learning neural audio embeddings for grounding semantics in auditory perception. *Journal of Artificial Intelligence Research*, 60, 1003–1030.
- Kisselew, M., Padó, S., Palmer, A., Šnajder, J. (2015, April). Obtaining a better understanding of distributional models of German derivational morphology. (pp. 58–63). London, UK: Association for Computational Linguistics.
- doi Landauer, T. K., & Dumais, S. T. (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. .
- doi Levy, O., Kenett, Y. N., Oxenberg, O., Castro, N., De Deyne, S., Vitevitch, M. S., Havlin, S. (2021). Unveiling the nature of interaction between semantics and phonology in lexical access based on multilayer networks. *Scientific Reports*, 11 (1), 1–14. .
- doi Manning, C., Surdeanu, M., Bauer, J., Finkel, J., Bethard, S., McClosky, D. (2014). The Stanford CoreNLP Natural Language Processing Toolkit. *Proceedings of 52nd annual meeting of the association for computational linguistics: System demonstrations* (pp. 55–60). Baltimore, Maryland: Association for Computational Linguistics.
- doi Marelli, M., Amenta, S., Crepaldi, D. (2015). Semantic Transparency in Free Stems: The Effect of Orthography-Semantics Consistency on Word Recognition. *Quarterly Journal of Experimental Psychology*, 68 (8), 1571–1583.
- doi Marelli, M., & Baroni, M. (2015). Affixation in semantic space: Modeling morpheme meanings with compositional distributional semantics. *Psychological Review*, 122 (3), 485–515.
- Mikolov, T. (2013, Jul 30). *word2vec*. Google Code Archive. Retrieved 2021-05-28, from <https://code.google.com/archive/p/word2vec/>
- Mikolov, T., Chen, K., Corrado, G., Dean, J. (2013). Efficient estimation of word representations in vector space. *1st International Conference on Learning Representations, ICLR 2013 – Workshop Track Proceedings*, 1–12. arXiv:1301.3781.
- doi Milin, P., Filipović Durdević, D., Moscoso del Prado Martín, F. (2009). The simultaneous effects of inflectional paradigms and classes on lexical recognition: Evidence from Serbian. *Journal of Memory and Language*, 60 (1), 50–64.
- doi Miller, G. A. (1995). WordNet: A lexical database for English. *Communications of the ACM*, 38 (11), 39–41.

- doi Monaghan, P., Shillcock, R. C., Christiansen, M. H., Kirby, S. (2014). How arbitrary is language? *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369 (1651).
- Moore, E. H. (1920). On the reciprocal of the general algebraic matrix. *Bulletin of the American Mathematical Society*, 26 (9), 394–395.
- doi Moscoso del Prado Martín, F., Kostić, A., Baayen, R. H. (2004). Putting the bits together: An information theoretical perspective on morphological processing. *Cognition*, 94, 1–18.
- doi Nikolaev, A., Chuang, Y., Baayen, R. H. (2022). A generating model for finnish nominal inflection using distributional semantics. *Accepted for publication in the Mental Lexicon*. Retrieved from <https://psyarxiv.com/ndtv2/download>.
- Ochshorn, R. M., & Hawkins, M. (2015). *Gentle: A robust yet lenient forced aligner built on kaldi*. (Available online at <http://lowerquality.com/gentle>).
- Park, J.A. (2013). *Spanish racoon cats*. Wikimedia Commons. Retrieved 2022-5-31, from https://commons.wikimedia.org/wiki/File:Spanish_racoon_cats.JPG (This file is licensed under the Creative Commons Attribution-Share Alike 3.0 Unported license).
- doi Plag, I., Homann, J., Kunter, G. (2017). Homophony and morphology: The acoustics of word-final S in English. *Journal of Linguistics*, 53 (1), 181–216.
- Polomé, E. C. (1967). *Swahili language handbook*. Washington, D.C.: Center for Applied Linguistics.
- Povey, D., Ghoshal, A., Boulianne, G., Burget, L., Glembek, O., Goel, N., ... Vesely, K. (2011, December). The kaldi speech recognition toolkit. *IEEE 2011 workshop on automatic speech recognition and understanding*. IEEE Signal Processing Society.
- doi Shafaei-Bajestan, E., Moradipour-Tari, M., Uhrig, P., Baayen, R. H. (2021). LDL-AURIS: a computational model, grounded in error-driven learning, for the comprehension of single spoken words. *Language, Cognition and Neuroscience*.
- Shafaei-Bajestan, E., Moradipour-Tari, M., Uhrig, P., Baayen, R. H. (2022). Semantic properties of english nominal pluralization: Insights from word embeddings. *arXiv*. arXiv:2203.15424.
- doi Shahmohammadi, H., Lensch, H. P.A., Baayen, R. H. (2021, November). Learning zero-shot multifaceted visually grounded word embeddings via multi-task training. *Proceedings of the 25th conference on computational natural language learning* (pp. 158–170). Online: Association for Computational Linguistics.
- Shillcock, R., Kirby, S., McDonald, S., Brew, C. (2001). Filled pauses and their status in the mental lexicon. *Proc. ITRW on Disfluency in Spontaneous Speech, Edinburgh, UK, 29–31 August 2001 (DiSS 2001)* (pp. 53–56). Edinburgh, UK: International Speech Communication Association.
- doi Siegelman, N., Rueckl, J. G., Lo, J. C. M., Kearns, D. M., Morris, R. D., Compton, D. L. (2022). Quantifying the regularities between orthography and semantics and their impact on group- and individual-level behavior. *Journal of Experimental Psychology: Learning Memory and Cognition*, 48 (6), 839–855.
- Sinclair, J. (1991). *Corpus, concordance, collocation*. Oxford: Oxford University Press.
- doi Tamariz, M. (2008). Exploring systematicity between phonological and context-cooccurrence representations of the mental lexicon. *The Mental Lexicon*, 3(2), 259–278. .
- doi Tomaschek, F., Plag, I., Ernestus, M., Baayen, R. H. (2019). Modeling the duration of word-final s in english with naive discriminative learning. *Journal of Linguistics*. (<https://psyarxiv.com/4bmwv>)

- Uhrig, P. (2018). Newsscape and the distributed little red hen lab – a digital infrastructure for the large-scale analysis of tv broadcasts. A.-J. Zwierlein, J. Petzold, K. Böhm, & M. Decker (Eds.), *Anglistentag 2017 in regensburg: Proceedings. proceedings of the conference of the german association of university teachers of english* (pp. 99–114). Trier: Wissenschaftlicher Verlag Trier.
- Uhrig, P. (2021). Large-Scale Multimodal Corpus Linguistics – The Big Data Turn (Habilitation thesis, unpublished manuscript). FAU Erlangen-Nürnberg.
- van der Maaten, L., & Hinton, G. (2008). Visualizing Data using t-SNE. *Journal of Machine Learning Research*, 9 (86), 2579–2605. Retrieved from <http://jmlr.org/papers/v9/vandermaaten08a.html>
- Vyagov, V. (2021). *Oranges (fruits)*. Wikimedia Commons. Retrieved 2022-5-31, from [https://commons.wikimedia.org/wiki/File:Oranges_\(fruits\).jpg](https://commons.wikimedia.org/wiki/File:Oranges_(fruits).jpg) (This file is licensed under the Creative Commons Attribution-Share Alike 4.0 International license.).
- Wang, B., Wang, A., Chen, F., Wang, Y., Kuo, C. C. J. (2019). Evaluating word embedding models: Methods and experimental results. *APSIPA Transactions on Signal and Information Processing*, 8 (1), e19.
- Yip, P.-C., & Rimmington, D. (2006). *Chinese: An essential grammar*. Routledge.

Address for correspondence

Elnaz Shafaei-Bajestan
 Eberhard Karls Universität Tübingen
 elnaz.shafaei@gmail.com
<https://orcid.org/0000-0002-1441-5608>

Co-author information

Peter Uhrig
 Friedrich-Alexander-Universität Erlangen-
 Nürnberg
 peter.uhrig@fau.de

R. Harald Baayen
 Eberhard Karls Universität Tübingen
 harald.baayen@uni-tuebingen.de

Publication history

Date received: 1 July 2022
 Date accepted: 14 March 2023
 Published online: 17 August 2023