On humans' (explicit) intuitions about the meaning of novel words

Article *in* Cognition · July 2024

DOI: 10.1016/j.cognition.2024.105882

4 authors:

Daniele Gatti
University of Pavia

Francesca Maria Rodio
University School for Advanced Studies IUSS Pavia

Luca Rinaldi
University of Pavia

Marco Marelli
Università degli Studi di Milano-Bicocca

On humans’ (explicit) intuitions about the meaning of novel words

Daniele Gatti[1]*, Francesca Rodio[1,2], Luca Rinaldi[1,3], & Marco Marelli[4,5]

[1] *Department of Brain and Behavioral Sciences, University of Pavia, Pavia, Italy*
[2] *Institute for Advanced Studies, IUSS, Pavia, Italy*
[3] *Cognitive Psychology Unit, IRCCS Mondino Foundation, Pavia, Italy*
[4] *Department of Psychology, University of Milano-Bicocca, Milano, Italy*
[5] *NeuroMI, Milan Center for Neuroscience, Milano, Italy*

Published on *Cognition*. Received 30 December 2023; Received in revised form 6 May 2024; Accepted 8 July 2024; Available online 17 July 2024.

https://doi.org/10.1016/j.cognition.2024.105882

*** Corresponding author:**

Daniele Gatti, Department of Brain and Behavioural Science, University of Pavia, Piazza Botta 6, 27100 Pavia, Italy. *e-mail address*: daniele.gatti@unipv.it

Author notes

Funding:

The contribution of MM was supported by the European Union (ERC-COG-2022, BraveNewWord, 101087053). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

The contribution of LR and FR was supported by funding from the BIAL Foundation (project: REMAP).

The contribution of LR was also supported by the Italian Ministry of Health (Ricerca Corrente 2023).

Competing interests:

Authors declare that they have no competing interests.

Open Practices Statement:

All data, scripts and codes used in the analysis are available at: https://osf.io/w234u/

The study was not preregistered.

Abstract

Pseudowords offer a unique opportunity to investigate how humans deal with new (verbal) information. Within this framework, previous studies have shown that, at the implicit level, humans exploit systematic associations in the form-meaning interface to process new information by relying on (sub-lexical) contents already mapped in semantic memory. However, whether speakers exploit such processes in explicit decisions about the meanings elicited by unfamiliar terms remains an open, important question. Here, we tested this by leveraging computational models that are able to induce semantic representations for out-of-vocabulary stimuli. Across two experiments, we demonstrate that participants' guesses about pseudoword meanings in a 2AFC task consistently align with the model's predictions. This indicates that humans' ability to extract meaningful knowledge from complex statistical patterns can affect explicit decisions.

Keywords

semantic memory, pseudowords, statistical learning, distributional semantic models

Introduction

Knoddled quocky ba boppi ziaowed tolque divords lurb: floal ribnier bureer. If the Introduction would start with this sentence nobody would (probably) understand the beginning of this article. The reason is simple: when looking for a meaning for these strings, an English speaker cannot find any direct connection with information stored in their semantic memory. *Knoddled* or *quocky* do (apparently) lack meaning, as they cannot be found in the vocabulary. However, while the vocabulary "knows" all the words, humans do not, and what might seem a meaningless string can be a word with a meaning that it is not (yet) known. This is indeed for example the case of low frequency words, like *lackadaisical*, that might not be known to a given speaker, but they might be able to activate intuitively a certain meaning. This makes lexical stimuli like *knoddled* or *quocky* intriguing from a scientific point of view: to all intents and purposes they are akin to unfamiliar existing words, allowing us to investigate if and how humans assign meaning to with novel (verbal and possibly meaningful) stimuli.

Verbal stimuli like *knoddled* or *quocky* are indeed generally labeled as "pseudowords", describing stimuli that are consistent with the phono- and orthotactical rules of a given language but are not attested in the lexicon of that language, and thus are not familiar to a given speaker. Contrary to the naïve perspective described above, in recent years, several studies have shown that semantic effects can be observed during the processing of these out-of-vocabulary stimuli (e.g., Bonandrini et al, 2023; Hendrix & Sun, 2021; Pugacheva & Günther, 2024; Sulpizio et al., 2021), that the same mechanisms governing word meaning can also subserve pseudowords processing (Gatti et al., 2023), and that humans are able to reliably assign affective content to these stimuli (e.g., Aryani et al., 2020; Gatti et al., 2024). These findings can be interpreted in terms of non-arbitrary components of language processing, like systematic form-to-meaning mapping (Dingemanse et al., 2015; for evidence on pseudowords processing see: Cassani et al., 2020; Chuang et al., 2021), that is humans' tendency to detect systematic and statistical regularities in the (language) environment (Romberg & Saffran, 2010; Vidal et al., 2021). In line with this view, previous studies have also shown that humans are able to exploit these mechanisms across a broad range of linguistic processes, in the grammatical, orthographical, phonological, and even semantic domains (for a review: Bogaerts et al., 2021; Christiansen, 2019).

The way humans attribute potential meaning to pseudowords has been recently investigated thanks to the methodological advancements in distributional semantic models (DSMs). Briefly, DSMs represent word meanings as high-dimensional numerical vectors induced from large corpora of natural language, under the assumption that the contexts in which words occur are informative of their meanings (Harris, 1954; Wittgenstein, 1953). Thus, words that are used in similar contexts in language will be located to nearby points in a semantic space, and the cosine of the angle between their vectors can be taken as a measure of how (semantically) related these words are (Günther et al., 2019; Mandera et al., 2017). Interestingly, it has been shown that it is even possible to induce representations for verbal stimuli that are not included in the training set (i.e., when they are out-of-vocabulary words) by modeling them as a sum of vectors representing the sequences of *n* contiguous letters (labeled as n-grams) composing it, that is by quantifying the distributional patterns of their sub-word information (Bojanowski et al., 2017). This approach has been used to estimate the "meaning" of pseudowords, that is the semantic pattern that an unfamiliar letter string can elicit in a speaker of the language (for a graphical representation see: Figure 1a, b, c, d).

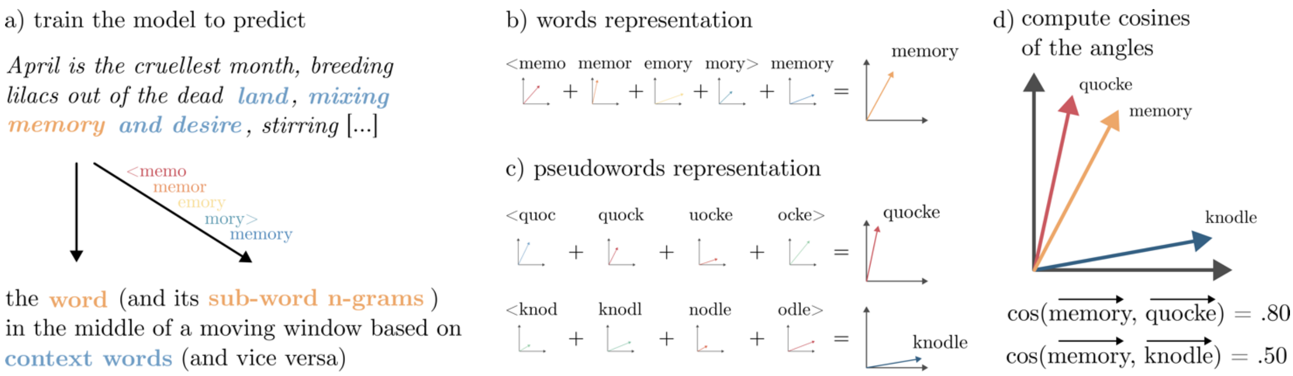

Figure 1. Schematical representation of how the DSM used in the present study is trained (a), represent the meaning of words (b) and pseudowords (c) and how to compare the similarity between the semantic pattern elicited by each string of letters (d).

By using this type of DSMs, it has been demonstrated that the semantic neighborhood density of pseudowords (i.e., how similar a pseudowords is to the five closest words in the semantic space) predicts humans' responses in lexical decision, with slower rejection latencies for pseudowords with denser semantic neighborhood (Hendrix & Sun, 2021; see also: Bonandrini et al., 2023) and that, in priming tasks, the more similar the meaning of a prime-word to the semantics elicited by a target-pseudoword, the slower participants' rejection latencies (Gatti et al., 2023). Overall, these studies showed that pseudowords can be indicative of meaning since – at least at the implicit level – humans' behavior is affected by the semantic pattern elicited by them. However, while it might seem intuitively reasonable that a speaker would exploit systematic sublexical patterns while processing new (verbal) information, to what extent these might affect explicit intuitions remains an open question. That is, explicit tasks focus on conscious, deliberate comparisons, potentially missing the subtle, automatic associations that influence decision-making at a chronometric level, as the ones employed by previous works. Observing such a reliance on systematic sublexical patterns while processing new (verbal) information also on a more explicit level would indicate that humans could be able to exploit linguistic statistical regularities in a "productive" way. That is, for example, this ability could emerge when there is the need to generate a novel label for a new concept, with possibly this process being optimizable for more applied reasons. Furthermore, in previous works modelling estimates were post-hoc obtained to describe previously collected behavioral data. That is, it was tested whether the model was able to account for existing phenomena; however, such approach missed a central contribution of computational psychology, namely prediction (e.g., Sun, 2008): a priori independently generating a certain output (the quantification of a given property, a set of automatically produced stimuli) that is then empirically tested in experiments involving human participants.

In the present study, we aimed to probe this possibility by conducting two-alternative-forced-choice (2AFC) experiments adopting a predictive approach. In both experiments, DSMs were applied to automatically produce stimulus sets in which a target was paired with two alternatives, the former maximally semantically related, the latter randomly associated. Participants were presented with the target, a string of letters corresponding to a word in Experiment 1 and a pseudoword in Experiment 2. They were then asked to indicate which of two alternatives (pseudowords in Experiment 1 and words in Experiment 2) was semantically more similar to the target stimulus. Notably, by showing (pseudo)words in isolation (i.e., without context) we aimed to directly test how sub-word components affect human behavior by removing possible contextual effects. We expected that, if humans are capable of assigning meaning to pseudowords at an explicit level, we would observe behavioral estimates aligned with the independently obtained model predictions.

Experiment 1

Methods

Participants

Sample size was determined a priori by means of a simulation procedure. We chose to include as coefficient for the effect size an extremely conservative value, b = .21 (i.e., probability = 55%). The choice to use this value was driven by the fact that we expected the possible semantic effect elicited by pseudowords to be small. The simulation showed that, using 50 experimental stimuli, the design employed here would have reached a power of 95% when including at least 55 participants (with an α = .05).

Sixty students participated in the study (7 males, *M* age = 21.75 years, *SD* = 2.55, age range = 19 – 34). All participants were native Italian speakers, had normal or corrected to normal vision and were naïve to the purpose of the study. Informed consent was obtained from all participants before the experiment. The protocol was approved by the psychological ethical committee of the University of Pavia and participants were treated in accordance with the Declaration of Helsinki.

Distributional semantic model

The DSM used here was *fastText* (Joulin et al., 2016), and in particular the Italian pre-trained vectors (Grave et al., 2018). We employed *fastText* because of its ability to compute semantic representations for both words and pseudowords. Indeed, *fastText* is based on the idea (originally proposed by Schütze, 1992; and realized computationally by Bojanowski et al., 2017) to take into account sub-word information and induce representations as the sum of the vectors of the letter n-grams associated with a given string. That is, *fastText* computes the semantic representation of a word as the sum of the vector of the full string plus all the vectors of the 5-grams that compose it.

A similar approach can be applied to unattested strings in order to try to capture the semantic information associated with pseudowords like *futmaw*. Of course, in this latter case, the induced representation will not consider the <futmaw> vector (since it does not exist by itself), but only the sum of its embedded n-grams. It should also be noted that, since *fastText* hashes n-grams into bins, including the ones that were not observed in the training data, out-of-vocabulary n-grams will be associated to a random vector (or to the vector corresponding to another n-gram in case of a collision). Note that this is not an issue when working with word-like pseudowords (as the ones employed here), as the overall number of out-of-vocabulary n-grams will be very low, thus limiting representational errors.

The model was trained on Common Crawl and Italian Wikipedia (around 11 billion words) using the Continuous Bag of Words (CBoW) method, an approach originally proposed by Mikolov and colleagues (2013), with 300 dimensions and a co-occurrence window of 5 words. When using CBoW, the obtained vector dimensions capture the extent to which a target element is reliably predicted by the linguistic contexts in which it appears, where "context" is represented as the words contained in a fixed size window around the target word. Specifically, the CBoW model will induce a representation for a given target w_0 based on context words w_{-n}, ..., w_{-1}, w_1, ..., w_n.

Using *fastText*, we therefore obtained semantic representations for the words and pseudowords included in the present Experiment. For each pair, we computed a semantic-relatedness index based on the cosine of the angle formed by vectors representing the meanings of the corresponding strings. The higher the cosine of the angle, the more semantically related the letter strings are expected to be, as estimated by the model.

Stimuli

Fifty triplets including one word and two pseudowords were automatically produced as stimuli. Firstly, using *fastText* (Joulin et al., 2016) on the Italian pre-trained vectors (Grave et al., 2018) we retrieved vector representations (see above the section **Distributional semantic model** for more information on *fastText*) for the 15,000 most frequent nouns and adjectives from the Italian SUBTLEX (http://crr.ugent.be/subtlex-it/) and for all the pseudowords included in Vergallito and colleagues (2020). Vergallito and colleagues (2020) report response latencies for 1,121 words and 1,121 pseudowords in a typical lexical decision task. Pseudowords were created using Wuggy (Keuleers & Brysbaert, 2010), a pseudoword generator that is able to create orthographic strings that respect the orthotactic rules of a given language (Italian, in the case of Vergallito et al., 2020).

Because *fastText* is based on very large natural language corpora and might have ended up including some non-existent string by mistake, we systematically checked whether a "whole-pseudoword" vector was available in the corpus for the pseudowords included. In such cases, indeed, *fastText* could learn distributional patterns about these pseudowords as if they were meaningful elements, even if their occurrence was based on errors and typos. A full vector representation was available for none of the considered pseudowords.

Then, after obtaining such semantic representations, we computed the cosine of the angle formed by each possible word-pseudoword pair vectors. The higher the cosine value, the more semantically related the letter strings are expected to be, as estimated by the model. Additionally, using the *stringdist R* package (Van der Loo, 2014) we computed the Levenshtein distance for each possible word-pseudoword pair. The Levenshtein distance measures the orthographic distance between two strings of characters by quantifying the minimum number of single-character edits (e.g., insertions, deletions, or substitutions) required to change one element into the other.

Finally, we defined the 50 triplets included as stimuli. Each triplet included one word and two pseudowords (e.g., as word: *zuffa*; as pseudowords: *umalo* and *tallarni*). Specifically, for each word, we retrieved two pseudowords: one was randomly selected among the ones with the closest vectors, while the other was equally random but semantically less similar. That is, among the pseudowords in each triplet, one was selected as related to the target word (i.e., it had a relatively high similarity index, with cosine similarity comprised between .20 and .44) and the other was selected as unrelated (i.e., it had a relatively low similarity index, with a cosine similarity comprised between .05 and .00007). The minimum cosine difference between the related and the unrelated pseudoword in a given triplet was .18. Among all the possible triplets, the 50 triplets eventually included were selected based on other linguistic indexes: length and Levenshtein distance. That is, in order to avoid that one of the two pseudowords was orthographically more similar to the target word or systematically longer/shorter than the other pseudoword, we selected the triplets that were more balanced across these indexes. This was tested by inspecting the histograms of the distributions and performing two two-samples Kolmogorov-Smirnov Test considering pseudowords lengths and Levenshtein distances, all Ds $< .12$, all ps $> .91$. Additionally, none of the pseudowords included was a pseudocompound (i.e., a concatenation of two existing words) or a novel derived form (i.e., a combination of existing stems and affixes), more specifically none of the pseudowords rhymed with the target word or shared with it the first letter, nor included recognizable suffixes, and no combination of evident prefix-like onsets (with length > 3) and existing words was included in our set of stimuli[1,2].

Procedure

Participants were tested using Psychopy (Pierce, 2007, 2009; Pierce & MacAskill, 2018; Pierce et al., 2019) through the online platform Pavlovia (https://pavlovia.org/).

Participants were told that they would have been presented with a word and two pseudowords (i.e., pronounceable out-of-vocabulary strings of letters). They were instructed that, although the pseudowords shown were unfamiliar, they could intuitively evoke a certain meaning, and that their task was to think about that potential meaning in order to judge which of the two pseudowords was more similar to the target word. They were also instructed to take all the time they needed for each trial.

Each trial started with a fixation cross (presented for 500 ms), then in the same screen a word was shown in uppercase letters in the upper part of the screen and two pseudowords in lowercase letters

[1] We discarded potential pseudowords based on suffixes only. The decision to consider only suffixes and not also prefixes was driven by the fact that it would have been extremely complex to exclude pseudowords based also on all the Italian prefixes (see: https://it.wikipedia.org/wiki/Prefissi_e_prefissoidi_della_lingua_italiana) as some of them, like the *a-* (indicative of negation) would have automatically excluded all the pseudowords beginning with *a-* (and a large number of Italian words start with *a-* but are not indicative of negation). However, note that no pseudoword having evident prefix-like onsets (**with length > 3**, e.g., anti-, contro-, extra-).

[2] For a complete list of Italian suffixes see: https://www.treccani.it/enciclopedia/suffissi_(La-grammatica-italiana)/

in left and right positions (until the participant's response). Participants indicated the chosen pseudoword with left and right keypresses (A and L). Participants' responses ended the trial and a blank screen (presented for 1000 ms) followed, then the next trial began (see Figure 2). Order of trials was counterbalanced across participants. The position on the screen (left vs. right) of related and unrelated words was counterbalanced (i.e., half of related words appeared on the left part of the screen and the other half on the right).

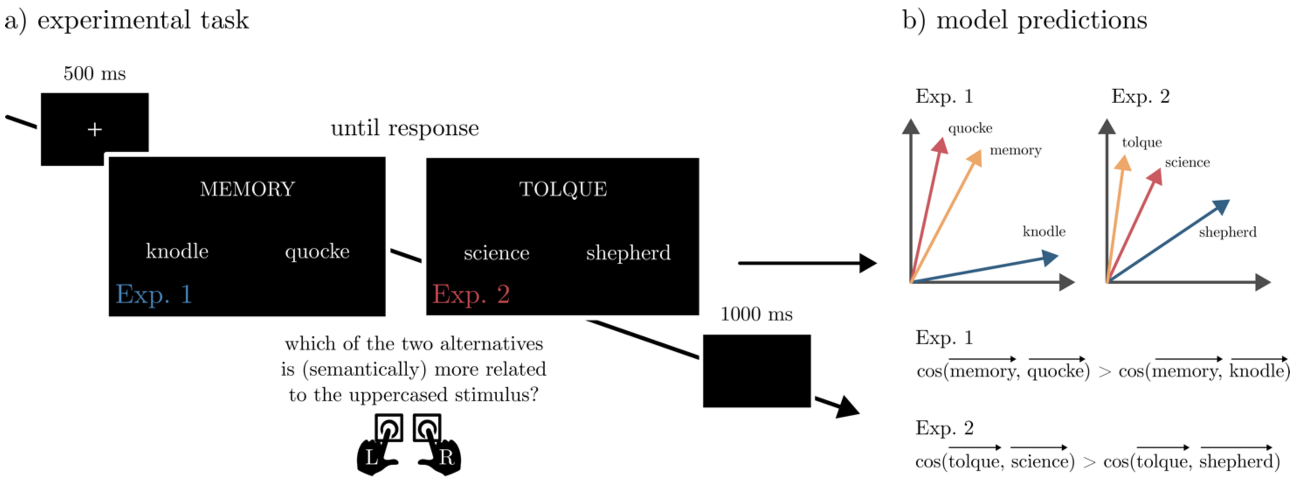

Figure 2. Schematical representation of the task used across Experiment 1 and Experiment 2; participants were shown one uppercased string of letters (a word in Experiment 1 and a pseudoword in Experiment 2) and were asked to indicate which one of the two lowercased strings of letters (pseudowords in Experiment 1 and words in Experiment 2) presented in the left and right of the screen was (semantically) more similar to the target word (note that for exposition the stimuli are in English, while the stimuli actually used were in Italian). They were also instructed that, although the pseudowords shown were out-of-vocabulary they could have been able to intuitively evoke a certain meaning, and that their task would have been to think at that potential meaning while solving the task (a). In each trial, one of the two alternatives (represented with the red vector) was predicted to be more related to the target stimulus (represented with the yellow vector) as compared with the other alternative according to the DMS used (b).

Data analysis and results

All the analyses were performed using *R*-Studio (RStudio Team, 2015). Data was analyzed through a mixed-effects approach, which incorporates both fixed-effects and random-effects (associated to participants and items) and allows for managing non-independency of the observations at both participants and item level (Baayen et al., 2008). Generalized linear mixed models (GLMMs) were run using the *lme4 R* package (Bates, et al., 2015) and were estimated on a binomial distribution.

The dependent variable was participants' binomial response (i.e., trials in which they selected the pseudoword that was produced by the model were coded as 1s, and trials in which they chose the other were scored as 0s). Hence, we tested whether participants selected as "related" the pseudoword produced by the DSM more frequently than the unrelated one. We first estimated a GLMM having participants' binomial responses as dependent variable and participants and items as random intercepts. That is, this model included only the intercept and random effects, allowing to test if participants' binomial responses differ from chance level. The estimates in GLMMs fitted on a binomial distribution are provided in log-odds (i.e., logit), thus if probability = .5, then the odds = 1 and the log-odds = 0. Thus, an estimate significantly higher than 0 indicates that participants' responses aligned towards the ones produced by the DSM. Indeed, this was the case: results indicated that participants reliably selected as "related" the pseudoword produced as such by the DSM, $z = 6.50$, $p < .001$, $b = .68$, prob. = 66%, *Pseudo-R^2 (total)* = .12 (Figure 3a).

Then, we tested whether participants' judgements could be predicted by the variables considered in the definition of the item set, namely: cosine similarity, orthographic length and Levenshtein distance. To do this, for each stimulus and for each variable, we computed the difference between the score for the related pseudoword vis-à-vis the unrelated one. For example, delta cosine similarity was computed by subtracting the cosine between the unrelated pseudoword and the target word from the cosine between the related pseudoword and the target word. Thus, we expect that participants' tendency to select the related pseudoword would increase at increased delta cosine similarity (i.e., as it should be easier to detect the "related" one). Broadly, these measures index how much the related pseudoword is more semantically or orthographically related to the target word, or longer as compared with the unrelated pseudoword. We thus estimated a GLMM having participants' binomial responses as dependent variable and participants and items as random intercepts. Delta cosine similarity, delta length, and delta Levenshtein distance were additively included as continuous predictors. Results are reported in Table 1 (*Pseudo-R^2 (total)* = .12, *Pseudo-R^2 (marginal)* = .01) and showed that delta cosine similarity predicted participants' performance, thus indicating that the higher the cosine similarity between the model-produced related pseudoword and the target word (as compared to the cosine similarity between the unrelated pseudoword and the target word), the higher the proportions of judgements aligned with the prediction of the model (Figure 3b). The other linguistic predictors were not significant.

Table 1. Results of the GLMM on participants' judgements including linguistic variables as predictors estimated in Experiment 1.

FIXED EFFECT	***b***	***z-value***	***p-value***
Δ cosine similarity	**4.11**	**2.05**	**.04**
Δ length	.09	1.90	.057
Δ Levenshtein distance	-.17	-1.92	.054

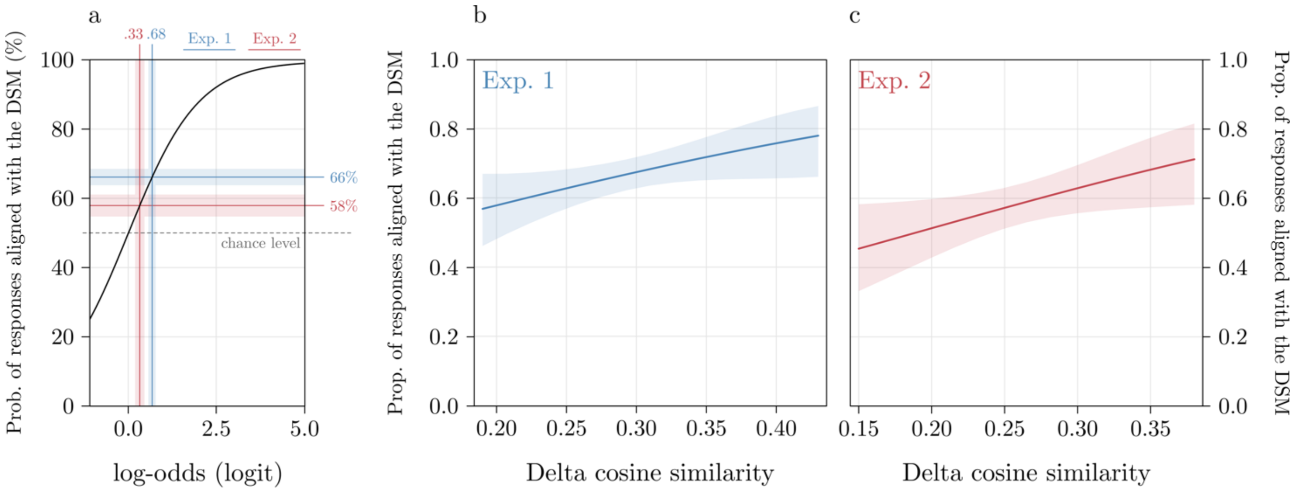

Figure 3. Plots illustrating the results of the GLMMs estimated in Experiment 1 (blue lines) and Experiment 2 (red lines) on participants' judgements showing that participants' classified as "related" the pseudoword considered as related by our DSM more frequently than the unrelated one (a); plots illustrating the results of the GLMMs including delta cosine similarity as predictor in Experiment 1 (b) and Experiment 2 (c). Across both experiments, the higher the delta, the higher participants' tendency to select as "related" the item predicted by the model.

Experiment 2

Methods

Participants

Power analysis was identical to Experiment 1. Sixty students participated in the study (22 males, *M* age = 25.2 years, *SD* = 3.45, age range = 19 – 35); none of them participated in Experiment 1. All participants were native Italian speakers, had normal or corrected to normal vision and were naïve to the purpose of the study. Informed consent was obtained from all participants before the experiment. The protocol was approved by the psychological ethical committee of the University of Pavia and participants were treated in accordance with the Declaration of Helsinki.

Distributional semantic model

The DSM used here was identical to Experiment 1.

Stimuli

The definition of the item set was similar to Experiment 1 and moved from the same pool of words and pseudowords but, in this case, the 50 triplets included one target pseudoword and two words as alternative. Among the words in each triplet, one was produced by the DSM as related to the target pseudoword (i.e., it had a high similarity index, with cosine similarity comprised between .20 and .44) and the other was estimated as unrelated (i.e., it had a low similarity index, with a cosine similarity comprised between .13 and .00003). The minimum cosine similarity difference between related and unrelated words was .14. The stimuli were balanced by inspecting the histograms of the distributions and performing a two-sample Kolmogorov-Smirnov Test considering word lengths, Levenshtein distance, frequencies of words and lemmas as retrieved from the Italian SUBTLEX (http://crr.ugent.be/subtlex-it/), all *D*s < .18, all *p*s > .39. Additionally, the two words included in each triplet were matched for gender (i.e., generally in Italian male and female words end with different letters), part of speech (i.e., noun or adjective) and number (i.e., singular or plural), and none of the words rhymed with the pseudoword nor shared with it the first letter. As for Experiment 1, none of the pseudowords included recognizable suffixes.

Procedure

The procedure was identical to Experiment 1; the only difference was that participants were instructed to judge which of the two words was more similar to the target pseudoword.

Data analysis and results

The first part of the data analysis was identical to Experiment 1. Results indicated that participants reliably selected as "related" the word produced as such by the DSM, $z = 2.45$, $p = .01$, $b = .33$, prob. = 58%, *Pseudo-*R^2 *(total)* = .20 (Figure 3a).

In the second part of the analyses, over and above delta cosine, delta length and delta Levenshtein distance, we also included the differences between the form and lemma frequency of the alternatives. These two new measures index how much the related word is more frequent as compared with the unrelated word. Results are reported in Table 2 (*Pseudo-*R^2 *(total)* = .20, *Pseudo-*R^2 *(marginal)* = .04) and showed that delta cosine similarity predicted participants' decision, thus indicating that the higher the cosine similarity between the model-produced related word and the target pseudoword (as compared to the cosine similarity between the unrelated word and the target pseudoword), the higher the proportions of judgements aligned with the prediction of the model (Figure 3c). The other linguistic predictors were not significant.

Table 2. Results of the GLMM on participants' judgements including linguistic variables as predictors estimated in Experiment 2.

FIXED EFFECT	*b*	*z-value*	*p-value*
Δ cosine similarity	**4.74**	**2.17**	**.03**
Δ length	.10	1.49	.14
Δ form frequency	.23	.47	.64
Δ lemma frequency	-.35	-.75	.45
Δ Levenshtein distance	-.15	-1.45	.15

Control analyses

In this section we present several control analyses performed to evaluate the specificity of the observed effects within the (systematic component of the) Italian language and to rule out possible trivial orthographic effects. On the one hand one might argue that the observed effect could be ascribed to language-independent effects related to similarity between linguistics sounds and their referents (e.g. phonosymbolic or iconic patterns, e.g., Lockwood et al., 2016; Sidhu et al., 2022), on the other hand one might argue that the observed effect could be traced back to simply the impact of orthographic neighbors (as in the turple effect; Forster & Hector, 2002) or of embedded strings (e.g., Bowers et al., 2005). This latter point is particularly relevant as excluding such impact would indicate that that the findings of the present study should be mainly traced back to the distributional history of the n-grams composing the (pseudo)words.

To exclude that the observed effect could ascribed to language-independent aspects, across both Experiment 1 and Experiment 2, we recoded participants' responses according to the estimates extracted from fastText DSMs pre-trained on Finnish and Basque (i.e., two languages that are typologically far from Italian). In fact, if the observed effect were due to general meaning impressions related to word forms (via iconicity or phonosymbolism, for example) we would find it even when using models trained on different languages than the one that is actually tested. Conversely, if the Italian-trained model outperformed the Finnish and Basque-trained ones, the observed effect should be interpreted as genuinely dependent on the distributional patterns at the sublexical level, and thus on the human ability to build upon it via statistical learning.

Models for a number of languages are available here: https://fasttext.cc/docs/en/crawl-vectors.html. We extracted from the Finnish and the Basque models vector representations for the words and the pseudowords included in both Experiment 1 and Experiment 2 and tested whether participants' responses aligned with their predictions.

Results showed that participants' responses did not align with the Finnish model, $z = .002$, $p = .99$, $b = .0003$, prob. = 50%, or the Basque model, $z = .89$, $p = .37$, $b = .12$, prob. = 53%. Similar patterns were found for Experiment 2 across both Finnish, $z = 1.48$, $p = .14$, $b = .20$, prob. = 55%, and Basque, $z = .77$, $p = .44$, $b = .11$, prob. = 52%[1]. These results indicate that the effects observed in Experiment 1 and Experiment 2 can be traced back to humans' sensitivity to meaningful patterns of letters which they were exposed during their (linguistic) experience, and not to the reliance on iconic or phonosymbolic cues in the adopted stimuli[3].

We hence implemented a way to estimate the semantic activation for pseudowords that did not rely on n-gram distributions, but rather on the impact of its orthographic neighbors. Therefore, as additional control across both Experiment 1 and Experiment 2, we recoded participants' responses according to the estimates of an approach extending the orthography-to-semantics (OSC) analyses proposed by Marelli and Amenta (2018), and that has seen a similar application to pseudowords in Hendrix and Sun (2021). In this case the vector of each pseudoword was defined as the average vector of its k closest orthographic neighbors (with $k = 5$ following Hendrix & Sun, 2021) among the 20k most frequent words attested in the Italian fastText model used (following Hendrix & Sun, 2021, and Gatti et al., 2023). Notably, to exclude the effect of sub-word information in the processes at hand, in this case the vector representations retrieved from fastText did not include sub-word information but only the "whole-word" vector of each stimulus (i.e., corresponding to a classical word2vec approach; Mikolov et al., 2013).

Overall, participants' responses significantly aligned with the prediction of the OSC model only in Experiment 1, $z = 3.20$, $p = .001$, $b = .40$, prob. = 60%, while in Experiment 2 this was not the case, $z = .86$, $p = .39$, $b = .12$, prob. = 53%. Notably, even though in Experiment 1 participants' responses

[3] Notably, across all these models, except for the Basque one in Experiment 2, we had to drop the intercept of the participants due to singular fit, thus indicating that this portion of the random effect did not contribute to explaining the observed variability.

aligned with the OSC model, the effect observed when the predictions included subword information (i.e., the one tested in Experiment 1; prob. = 66%) was substantially larger.

These control analyses indicate that the effects observed are related to actual systematicity in the distributional history of sublexical units within a given language, and not trivially dependent on the processing of orthographic neighbors, thus toning down interpretations related to phenomena such as the turple effect or the impact of embedded strings (e.g., Bowers et al., 2005; Forster & Hector, 2002).

Discussion

In the present study, we investigated whether humans' intuitions about the meaning of novel words can be observed in a task requiring an explicit decision between two alternatives. We took advantage of a distributional semantic model (DSM), namely *fastText* (e.g., Bojanowski et al., 2017), able to approximate the semantic information evoked by sub-word units in language and explored whether such prediction is aligned with participants' intuitions. Across two 2AFC experiments, participants were shown a target item (i.e., a word in Experiment 1 and a pseudoword in Experiment 2) and two other alternative items (i.e., two pseudowords in Experiment 1 and two words in Experiment 2) and were then asked to indicate which of two strings of letters they felt to be more related to the target item, in terms of the meanings they evoked in their semantic memory. Results showed that participants reliably selected the stimulus produced by the DSM used. Notably, the language used in the present study – Italian – is completely transparent and thus, by controlling for orthographic information, we also directly ruled out possible phonological constraints. Follow-up analyses further revealed that the higher the difference in cosine similarity between the alternatives and the target the higher the probability of responding consistently with the model predictions, thus ruling out the possibility that the effect was item-dependent. That is, while the fact that participants responded consistent with the prediction of the model higher than chance could have been caused by the structure of the item set (i.e., the distribution of the population could be centered on chance level and we could have sampled more items from one of the two halves), the follow-up analysis directly traces back the observed effect to a specific process (i.e., the semantic information evoked by sub-word units as predicted by the DSM).

Overall, these findings extend previous evidence suggesting that, at the implicit level, humans are sensitive to the semantic patterns elicited by novel words: this effect can be indeed observed also when requesting explicit intuitions. Overall, the present study indicates that humans can exploit distributional information in their language to explicitly make sense of novel (seemingly meaningless) stimuli. Because the task we employed is very simple, this also speaks in favor of the generalizability of this phenomenon to many everyday life situations. Moreover, whereas previous studies post-hoc produced model estimates for stimuli employed in existing experiments, here we applied the model to directly and automatically generate an item set that was then administered to participants. In this perspective, the present study speaks for the reliability of the *fastText* estimates and their wide applicability in cognitive research; the model does not only produce robust measures but can also predict novel unexplored phenomena.

The effects described can be traced back to humans' tendency to detect systematic and statistical regularities in the (language) environment (Romberg & Saffran, 2010; Vidal et al., 2021) and thus can be framed within non-arbitrary perspectives on language, with specific reference to systematic mapping (Dingemanse et al., 2015). Within this context, systematic mapping refers to the reliable statistical relationships between sub-lexical structures and semantic features (e.g., Nölle et al., 2018). Reliance on systematic (but also iconic) pattern has been shown within early-age word learning (Imai et al., 2008; Monaghan et al., 2011; Monaghan & Roberts, 2021; Nielsen & Dingemanse, 2021) and, more generally, in scaffolding the production and comprehension of language (Perry et al., 2018). Consistent with this theoretical line, several studies have shown that, when processing novel words or words from an unfamiliar language, speakers exploit form-related cues to activate meaning information (Cassani et al., 2020; Forster & Hector, 2002; Louwerse & Qu, 2017). More specifically, our results extend those described by previous studies employing (linear) discriminative learning (i.e., implementing linear mappings between pseudowords form and semantic vectors; Baayen et al., 2019) algorithms to account for the systematic relation between pseudowords forms and their meanings (e.g., Cassani et al., 2020; Chuang et al., 2021). In parallel, seminal (distributed) connectionists approaches argued that morphology (i.e., how words are formed) might reflect a learned sensitivity to the systematic relationships among the surface forms of words and their meanings (e.g., Plaut & Gonnerman, 2000). Consistent with this, recently Ulicheva and colleagues (2020) have shown across multiple tasks that (in English) suffixes are highly informative of parts of speech, and that readers are sensible to this sub-lexical regularity. Critically, in explaining humans' behavior as observed in the present study, we may argue that *fastText*'s architecture, by relying on sub-word information, might computationally index the same form-to-meaning components that were algorithmically described by these pioneering approaches. Notably, strengthening the observed effects, while the Plaut & Gonnerman (2000) model required the modeller's input in the pre-segmentation phase (see: Rastle & Davis, 2008, but also: Stevens & Plaut, 2022), *fastText* works blindly, with this (possibly) providing new (and entirely bottom-up) ways to discovering the meanings of morphemes.

Building upon these theoretical perspectives and pushing the argument further, humans' ability to overtly exploit sublexical information when making sense of novel but plausible linguistic information can be linked to psychological models that describe lexical effects as an epiphenomenon of stable statistical patterns between form and meaning (e.g., Baayen et al., 2011, 2019; Harm & Seidenberg, 2004; Milin et al., 2017; Seidenberg, 1995). More specifically, in the model proposed by Harm and Seidenberg's (2004), semantics, orthography and phonology constitute the angles of a triangle and its sides represent the interconnected (and possibly bidirectional) nature of the processes at hand. Interestingly, in this model, semantics emerges following information running through both the ortho-semantic and the phono-semantic pathways. Considering this, we can interpret the effect of sublexical information when understanding novel words via the processing of orthographic information to activate meaning (i.e., the side of the triangle linking orthography and semantics) and thus ultimately describing form-to-meaning mapping.

Other relevant theoretical approaches can be used to explain our findings, like the dual route cascade model (DRC, see Coltheart et al., 2001) and the dual-route approach to orthographic processing (Grainger & Ziegler, 2011). According to the former model, humans would rely on two distinct pathways when recognizing a word: a lexical one, involved in recognizing words directly via representations of word forms stored in memory, and a non-lexical one, based on sublexical regularities and commonly used to decode unfamiliar letter strings. Given the absence of a lexical entry (i.e., the pseudoword does not exist in the vocabulary, and thus the lexical route cannot be activated), we can interpret our findings as humans' reliance on the non-lexical route and thus to the emergence of semantic information encoded at the n-gram level through the use of sublexical (orthographic) regularities. However, while in the classical DRC (e.g., Coltheart et al., 2001) lexical and sublexical pathways run parallelly, and thus the latter one cannot affect semantic processes, our results speak for possible interactions between the two routes. Alternatively, these results could be framed within the classical DRC by keeping the notion of parallel pathways, and by arguing that the lexical pathway could be sensitive to (distributionally) salient strings of letters, with this ability serving as a bridge to semantics. Finally, according to Grainger and Ziegler (2011), humans are thought optimize the mapping of form to meaning by using two different prelexical orthographic codes: a coarse-grained one which facilitates the access to word meaning by relying on the identification of highly informative letters, and a fine-grained one which is characterized by the detection of (pre-existing) relevant sublexical combinations of letters. We can interpret the results of the present study as a reader's reliance on the latter code, that can be activated (even) when the stimulus has no place on the vocabulary of a given language and explicitly exploited to make sense of novel (linguistic) information.

The observed effects can be further framed by drawing a parallel with episodic memory. When experiencing a new event, individuals can try to make sense of it by exploiting information from events they encoded during their lives (e.g., Tulving, 1993; 2002). This can be done at the declarative (i.e., explicit) level by navigating at will within the information stored in memory and then by retrieving it. Similarly, here, individuals are shown to (overtly) navigating within their semantic memory to search for a (possible) way to interpret the meaning of novel words. This encompasses the idea that semantic memory is a generative system that constantly deals with novel information, as supported by studies estimating that an adult speaker learns from 1.7 (Nation, 2006) to 11 (Nagy & Anderson, 1984) novel words per day (Brysbaert et al., 2016). In doing so, humans would take advantage of low-level featural elements found in the environment that, in the case of verbal stimuli, are ultimately quantifying the (distributional) learning history of sub-word units in language. Indeed, given the nature of the DSM applied here, these findings are consistent with a view of semantic memory as a cognitive system that taps onto general-purpose associative learning mechanisms (Günther et al., 2019). Pushing forward these generative capabilities of semantic memory, this explicit effect indicates that, in principle, humans could be able to exploit systematic regularities of sublexical units in a given language not only when processing novel words, but also when asked to generate novel labels for new (or existing) concepts. This topic can be of great applicative interest, and we believe constitutes the major future direction for the findings reported here.

Regarding other future directions, the method adopted here could be easily applied to answer novel empirical questions. For example, one might test to what extent bilingual individuals rely on L1 or on L2 when performing a similar task. Additionally, in the present study we used plausible linguistic stimuli shown in isolation; a possible extension is hence related to the use of pseudowords in context and/or stimuli that do not follow a given language's orthotactics. This allows to test how the reliance on sub-word information when dealing with novel linguistic stimuli can be generalized across different scenarios and tasks. This latter perspective is particularly intriguing as it would allow to clarify whether the findings of the present study, that is that (pseudo)word meaning can be extracted from the distributional history of the n-grams composing the stimulus, is dependent on the readability of the stimuli. Indeed, non-readable stimuli typically include sublexical elements that are extremely rare, if at all attested (e.g., "klvmst" or "rptglf"). Their associated distributions might hence not be informative enough to elicit any semantic intuitions, shaping readability as a crucial condition to set off semantic access. Interestingly, as *fastText* allows to specify the length of the sequences of letters (i.e., the n-grams) to be considered in the training phase, future studies could address this point by training ad-hoc models including information from shorter n-grams generally not considered in pre-trained models (e.g., uni-grams, bi-grams) and using the experimental procedure adopted here.

A possible source of concern is related to the use of computational modeling in predicting human behavior. On the one hand, indeed, one should be careful in inferring that model's parameters and algorithms can be directly applied to human cognition. On the other hand, the empirical evidence presented here indicates that there is a certain degree of overlap between *fastText* predictions and humans processing of novel information. When reasoning about *fastText* characteristics, one should always keep in mind that it is a resource build within natural language processing contexts with the explicit applicative scope of facilitating and improving text representation (and not pseudoword representation!). We believe that this latter point does strengthen even more our results as *fastText* architecture was not explicitly tuned for the material and processes we are investigating, but nevertheless it can be used to (successfully) capture humans' responses to such stimuli.

In conclusion, using DSMs we provide evidence that humans are able to exploit sub-word information when dealing with novel words in an explicit task, thus demonstrating that semantic (explicit) intuitions on the meaning of novel (verbal) stimuli can be traced back to domain-general associative mechanisms. Our findings directly support theories on the non-arbitrariness of language and provide novel insights into the distributed structure of human semantic memory.

References

Adriaans, F., & Kager, R. (2010). Adding generalization to statistical learning: The induction of phonotactics from continuous speech. *Journal of Memory and Language*, 62(3), 311–331.

Altmann, G., Dienes, Z., & Goode, A. (1995). Modality independence of implicitly learned grammatical knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*(4), 899.

Aryani, A., Isbilen, E. S., & Christiansen, M. H. (2020). Affective arousal links sound to meaning. *Psychological Science*, *31*(8), 978-986.

Baayen R. H., Chuang Y. Y., Shafaei-Bajestan E., Blevins J. P. (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*, 2019, Article 4895891.

Baayen R. H., Milin P., Đurđević D. F., Hendrix P., Marelli M. (2011). An amorphous model for morphological processing in visual comprehension based on naive discriminative learning. *Psychological Review*, 118(3), 438–481.

Baayen, R. H., Chuang, Y. Y., Shafaei-Bajestan, E., & Blevins, J. P. (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*, 2019.

Baayen, R. H., Davidson, D. J., & Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, *59*(4), 390-412.

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1), 1–48.

Bogaerts, L., Siegelman, N., & Frost, R. (2021). Statistical learning and language impairments: Toward more precise theoretical accounts. *Perspectives on Psychological Science*, 16(2), 319–337

Bonandrini, R., Amenta, S., Sulpizio, S., Tettamanti, M., Mazzucchelli, A., & Marelli, M. (2023). Form to meaning mapping and the impact of explicit morpheme combination in novel word processing. *Cognitive Psychology*, *145*, 101594.

Bojanowski, P., Grave, E., Joulin, A., & Mikolov, T. (2017). Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5, 135–146

Bowers, J. S., Davis, C. J., & Hanley, D. A. (2005). Automatic semantic activation of embedded words: Is there a "hat" in "that"?. *Journal of Memory and Language*, *52*(1), 131-143.

Brysbaert, M., Stevens, M., Mandera, P., & Keuleers, E. (2016). How many words do we know? Practical estimates of vocabulary size dependent on word definition, the degree of language input and the participant's age. *Frontiers in Psychology*, 7

Cassani, G., Chuang, Y. Y., & Baayen, R. H. (2020). On the semantics of nonwords and their lexical category. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *46*(4), 621.

Christiansen, M. H. (2019). Implicit statistical learning: A tale of two literatures. *Topics in Cognitive Science*, 11(3), 468–481.

Chuang, Y. Y., Vollmer, M. L., Shafaei-Bajestan, E., Gahl, S., Hendrix, P., & Baayen, R. H. (2021). The processing of pseudoword form and meaning in production and comprehension: A computational modeling approach using linear discriminative learning. *Behavior Research Methods*, *53*, 945-976.

Coltheart, M., Rastle, K., Perry, C., Langdon, R., & Ziegler, J. (2001). DRC: a dual route cascaded model of visual word recognition and reading aloud. *Psychological Review*, *108*(1), 204.

Dingemanse, M., Blasi, D. E., Lupyan, G., Christiansen, M. H., & Monaghan, P. (2015). Arbitrariness, iconicity, and systematicity in language. *Trends in Cognitive Sciences*, 19(10), 603–615

Forster, K. I., & Hector, J. (2002). Cascaded versus noncascaded models of lexical and semantic processing: The turple effect. *Memory & Cognition*, 30(7), 1106-1117.

Gatti, D., Marelli, M., & Rinaldi, L. (2023). Out-of-vocabulary but not meaningless: Evidence for semantic-priming effects in pseudoword processing. *Journal of Experimental Psychology: General*, 152(3), 851.

Gatti, D., Raveling, L., Petrenco, A., & Günther, F. (2024). Valence without meaning: investigating form and semantic components in pseudowords valence. *Psychonomic Bulletin & Review*, 1-13.

Grainger, J., & Ziegler, J. C. (2011). A dual-route approach to orthographic processing. *Frontiers in Psychology*, *2*, 54.

Grave, E., Bojanowski, P., Gupta, P., Joulin, A., & Mikolov, T. (2018). Learning word vectors for 157 languages. *arXiv Preprints*.

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.

Harm M. W., Seidenberg M. S. (2004). Computing the meanings of words in reading: Cooperative division of labor between visual and phonological processes. *Psychological Review*, 111, 662–720.

Harris, Z. S. (1954). Distributional structure. *Word*, 10(2–3), 146–162.

Hendrix, P., & Sun, C. C. (2021). A word or two about nonwords: Frequency, semantic neighborhood density, and orthography-to-semantics consistency effects for nonwords in the lexical decision task. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 47(1), 157–183.

Imai, M., Kita, S., Nagumo, M., & Okada, H. (2008). Sound symbolism facilitates early verb learning. *Cognition*, 109(1), 54-65.

Joulin, A., Grave, E., Bojanowski, P., Douze, M., Jégou, H., & Mikolov, T. (2016). Fasttext. zip: Compressing text classification models. *arXiv Preprints*.

Keuleers, E., & Brysbaert, M. (2010). Wuggy: A multilingual pseudoword generator. *Behavior Research Methods*, 42(3), 627–633.

Keuleers, E., Lacey, P., Rastle, K., & Brysbaert, M. (2012). The British Lexicon Project: Lexical decision data for 28,730 monosyllabic and disyllabic English words. *Behavior Research Methods, 44*(1), 287–304.

Lockwood, G., Dingemanse, M., & Hagoort, P. (2016). Sound-symbolism boosts novel word learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *42*(8), 1274.

Louwerse, M., & Qu, Z. (2017). Estimating valence from the sound of a word: Computational, experimental, and cross-linguistic evidence. *Psychonomic Bulletin & Review, 24(3),* 849-855.

MacAskill, M., & Peirce, J. (2018). *Building Experiments in PsychoPy*. Building Experiments in PsychoPy, 1-312.

Mandera, P., Keuleers, E., & Brysbaert, M. (2017). Explaining human performance in psycholinguistic tasks with models of semantic similarity based on prediction and counting: A review and empirical validation. *Journal of Memory and Language*, 92, 57–78.

Marelli, M., & Amenta, S. (2018). A database of orthography-semantics consistency (OSC) estimates for 15,017 English words. *Behavior Research Methods*, *50*, 1482-1495.

Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. *arXiv Preprints*.

Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.

Milin P., Feldman L. B., Ramscar M., Hendrix P., Baayen R. H. (2017b). Discrimination in lexical decision. *PLOS ONE*, 12, Article e0171935.

Monaghan, P., & Roberts, S. G. (2021). Iconicity and diachronic language change. *Cognitive Science*, 45(4), e12968.

Monaghan, P., Christiansen, M. H., & Fitneva, S. A. (2011). The arbitrariness of the sign: Learning advantages from the structure of the vocabulary. *Journal of Experimental Psychology: General*, 140(3), 325-347.

Nagy, W. E., & Anderson, R. C. (1984). How many words are there in printed school English? *Reading Research Quarterly*, 19(3), 304–330.

Nation, I. (2006). How large a vocabulary is needed for reading and listening? Canadian Modern Language Review, 63(1), 59–82. on *Psychological Science*, 16(2), 319–337

Nielsen, A. K., & Dingemanse, M. (2021). Iconicity in word learning and beyond: A critical review. *Language and Speech*, 64(1), 52-72.

Nölle, J., Staib, M., Fusaroli, R., & Tylén, K. (2018). The emergence of systematicity: How environmental and communicative factors shape a novel communication system. *Cognition*, 181, 93-104.

Peirce J, Gray JR, Simpson S, MacAskill M, Höchenberger R, Sogo H, Kastman E, Lindeløv JK. (2019) PsychoPy2: Experiments in behavior made easy. *Behavior Research Methods*. 51(1):195-203.

Peirce, J. W. (2007). PsychoPy—psychophysics software in Python. *Journal of Neuroscience Methods*, 162(1-2), 8-13.

Peirce, J. W. (2009). Generating stimuli for neuroscience using PsychoPy. *Frontiers in Neuroinformatics*, 2,343.

Perry, L.K., Perlman, M., Winter, B., Massaro, D.W., & Lupyan, G. (2018). Iconicity in the speech of children and adults. *Developmental Science*, 21(3), e12572.

Plaut, D. C., & Gonnerman, L. M. (2000). Are non-semantic morphological effects incompatible with a distributed connectionist approach to lexical processing?. *Language and Cognitive Processes*, *15*(4-5), 445-485.

Pugacheva, V., & Guenther, F. (20243). Lexical choice and word formation in a taboo game paradigm. *PsyArXiv preprintJournal of Memory and Language.*

Rastle, K., & Davis, M. H. (2008). Morphological decomposition based on the analysis of orthography. *Language and Cognitive Processes*, *23*(7-8), 942-971.

Romberg, A. R., & Saffran, J. R. (2010). Statistical learning and language acquisition. *Wiley Interdisciplinary Reviews: Cognitive Science*, 1(6), 906–914.

RStudio Team. (2015). RStudio: Integrated Development for R (Version 0.98.1074) [Computer software]. RStudio, Inc. http://www.rstudio.com/

Saffran, J. R., & Wilson, D. P. (2003). From syllables to syntax: Multilevel statistical learning by 12-month-old infants. *Infancy*, 4(2), 273–284.

Schutze, H. (1992). Word space. *Advances in Neural Information Processing Systems*, 5, 895–902.

Seidenberg M. S. (1995). Visual word recognition. In Miller J. L., Eimas P. D. (Eds.), *Handbook of perception & cognition: Vol. 11. Speech, Language & Communication* (pp. 137–179). Academic Press.

Sidhu, D. M., Vigliocco, G., & Pexman, P. M. (2022). Higher order factors of sound symbolism. *Journal of Memory and Language*, *125*, 104323.

Stevens, P., & Plaut, D. C. (2022). From decomposition to distributed theories of morphological processing in reading. *Psychonomic Bulletin & Review*, *29*(5), 1673-1702.

Sulpizio, S., Pennucci, E., & Job, R. (2021). The impact of emotional content on pseudoword recognition. *Psychological Research*, *85*, 2980-2996.

Sun, R. (2008). Introduction to computational cognitive modeling. *Cambridge handbook of computational psychology*, 3-19.

Thompson, S. P., & Newport, E. L. (2007). Statistical learning of syntax: The role of transitional probability. *Language Learning and Development*, 3(1), 1–42.

Tulving, E. (1993). What is episodic memory?. *Current Directions in Psychological Science*, 2(3), 67-70.

Tulving, E. (2002). Episodic memory: From mind to brain. *Annual review of psychology*, 53(1), 1-25.

Ulicheva, A., Harvey, H., Aronoff, M., & Rastle, K. (2020). Skilled readers' sensitivity to meaningful regularities in English writing. *Cognition*, *195*, 103810.

Ulicheva, A., Marelli, M., & Rastle, K. (2021). Sensitivity to meaningful regularities acquired through experience. *Morphology*, *31*, 275-296.

Van der Loo, M. P. J. (2014). The stringdist package for approximate string matching. *The R Journal*, 6(1), 111–122.

Vergallito, A., Petilli, M. A., & Marelli, M. (2020). Perceptual modality norms for 1,121 Italian words: A comparison with concreteness and imageability scores and an analysis of their impact in word processing tasks. *Behavior Research Methods*, *52*(4), 1599-1616.

Vidal, Y., Viviani, E., Zoccolan, D., & Crepaldi, D. (2021). A general-purpose mechanism of visual feature association in visual word identification and beyond. *Current Biology*, 31(6), 1261–1267.e3

Wittgenstein, L. (1953). *Philosphical investigations*. New York. MacMillan.