©2022, American Psychological Association. This paper is not the copy of record and may not exactly replicate the final, authoritative version of the article. Please do not copy or cite without authors' permission. The final article will be available, upon publication, on the *Journal of Experimental Psychology: General* via its DOI: 10.1037/xge0001304

Out-of-vocabulary but not meaningless: Evidence for semantic-priming effects in pseudoword processing

Daniele Gatti¹, Marco Marelli^{2,3} & Luca Rinaldi^{1,4}

¹ *Department of Brain and Behavioral Sciences, University of Pavia, Pavia (Italy)*

² *Department of Psychology, University of Milano-Bicocca, Milano, Italy*

³ *NeuroMI, Milan Center for Neuroscience, Milano, Italy*

⁴ *Cognitive Psychology Unit, IRCCS Mondino Foundation, Pavia (Italy)*

Author note

Correspondence concerning this article should be addressed to Daniele Gatti, Department of Brain and Behavioural Science, University of Pavia, Piazza Botta 6, 27100 Pavia, Italy. *e-mail address:* daniele.gatti@unipv.it

Funding: Italian Ministry of Health, grant Ricerca Corrente 2021 (LR)

Prior dissemination: A previous version of this article was uploaded on PsyArXiv (<https://psyarxiv.com/y2fm3/>)

Competing interests: Authors declare that they have no competing interests.

Data availability: All data, scripts and codes used in the analysis are available at: <https://osf.io/pt97n/>

Abstract

Non-arbitrary phenomena in language, such as systematic association in the form-meaning interface, have been widely reported in the literature. Exploiting such systematic associations previous studies have demonstrated that pseudowords can be indicative of meaning. However, whether semantic activation from words and pseudowords is supported by the very same processes, activating a common semantic memory system, is currently not known. Here, we take advantage of recent progresses from computational linguistics models allowing to induce meaning representations for out-of-vocabulary strings of letters via domain-general associative-learning mechanisms applied to natural language. We combined these models with data from priming tasks, in which participants are showed two strings of letters presented sequentially one after the other and are then asked to indicate if the latter is a word or a pseudoword. In Experiment 1 we re-analyzed the data of the largest behavioral database on semantic priming, while in Experiment 2 we ran an independent replication on a new language, Italian, controlling for a series of possible confounds. Results were consistent across the two experiments and showed that the prime-word meaning interferes with the semantic pattern elicited by the target pseudoword (i.e., at increasing estimated semantic relatedness between prime word and target pseudoword, participants' reaction times increased and accuracy decreased). These findings indicate that the same associative mechanisms governing word meaning also subserve the processing of pseudowords, suggesting in turn that human semantic memory can be conceived as a distributional system that builds upon a general-purpose capacity of extracting knowledge from complex statistical patterns.

Keywords: semantic priming, pseudowords, statistical learning, distributional semantic model, semantic memory.

Introduction

Language is a system of communication that uses symbols to convey deep and often complex meanings. These symbols are words, whose forms have no apparent a-priori connection with their meanings: the mapping between word form (e.g., *cat*) and its associated meaning (i.e., small four-legged mammal from the family of *Felidae*) is purely arbitrary and thus dependent mainly (or only) on cultural conventions. This theoretical view on meaning has been formally proposed at the beginning of the XX Century (e.g., De Saussure, 1916), but its assumptions can be traced back to Plato's writings, especially to the *Cratylus*.

Whether the relationship between word form and meaning is purely arbitrary or not has been at the center of an intense dispute for decades, which is still nowadays igniting a fervent debate (for a discussion see: Dingemanse et al., 2015). Against this conventional perspective, indeed, it has been argued that the principle of arbitrariness would account only partly for form-to-meaning correspondences, with recent research providing increasing evidence for two common types of non-arbitrary relationships in language: namely, iconicity and systematicity (Dingemanse et al., 2015). Iconicity, also referred to as sound-symbolism (e.g., Köhler, 1929; Sapir, 1929), refers to the resemblance-based mapping between properties of linguistic form and meaning: that is, properties of linguistic form can resemble their referent in some ways, as in the case of onomatopoeic words or prosodic modulation during speech (e.g., a speaker saying *taaaall* for emphasizing the height of a particularly tall skyscraper; Perniss & Vigliocco, 2014). Beyond iconicity, systematic mapping in the form-meaning interface have been widely reported, referring to reliable statistical relationships between sublexical structures and semantic features (e.g., Nölle et al., 2018).

Both iconic and systematic patterns are fundamental for early-age word learning (Imai et al., 2008; Monaghan et al., 2011; Monaghan & Roberts, 2021; Nielsen & Dingemanse, 2021), scaffolding the production and comprehension of language (Perry et al., 2018). From early after birth, humans are indeed particularly good at detecting statistical regularities in the (language) environment (Romberg & Saffran, 2010), with this ability likely being crucial also during adulthood and for the whole lifespan, especially considering that language users constantly learn new words (Brysbaert et al., 2016). Language use is indeed fundamentally contextual and probabilistic, with linguistic context exerting a pervasive influence on the form and content of human communication (Ramscar, 2019). In line with this possibility, supporting the idea that systematic regularities in language can be indicative of meaning, several studies have shown that, when processing new 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). In principle, therefore, assigning

meaning to new words could be driven by nuanced systematic form-to-meaning cues (Cassani et al., 2020), that are not necessarily characterized in linguistic terms (e.g., sublexical units like morphemes) on an explicit level. This is also indirectly supported by recent evidence demonstrating that, in a lexical decision task, the rejection times for pseudowords are modulated by their semantic neighborhood density, that is the average distance between the target stimulus and its 5 closest semantic neighbors, as estimated via distributional semantic models (DSMs) (Hendrix & Sun, 2021). DSMs are indeed particularly well suited to quantitatively characterize meaning, as they are conceived as computationally implemented theories of human semantic representations, building on word distributions in language experience (Günther et al., 2019; Landauer & Dumais, 1997).

In particular, DSMs represent word meanings as high-dimensional numerical vectors, induced from large amounts of natural language data: the assumption at the core of these models is the distributional hypothesis, according to which contexts in which words occur can approximate their meanings (Harris, 1954). Hence, words with similar meanings (i.e., with a similar distributional pattern across linguistic contexts) are mapped to nearby points in a semantic space (Günther et al., 2019). Critically, DSMs have been shown to be impressively high-performing across a wide range of psychological tasks such as multiple-choice tests (Bullinaria & Levy, 2012), word categorization (Baroni & Lenci, 2010), word relatedness ratings (Bruni et al., 2014), word naming and lexical decision (Marelli & Amenta, 2018), semantic priming (Günther et al., 2016), recognition memory (Gatti et al., 2021), thoroughly accounting for human performance (for a review see: Mander et al., 2017). Importantly, recent progress in DSMs provide us with a convenient way to induce representations for sub-word information: *fastText* (Bojanowski et al., 2017) is, in fact, a DSM that can quantify the distributional pattern of sub-word information in natural language. The benefit of this approach is that each verbal stimulus (e.g., a word) is modeled as a sum of vectors representing the sequences of n contiguous letters (labeled as n -grams) composing it: as such, in principle *fastText* can also be used to approximate the “meaning” of pseudowords, that is the semantic patterns that an unfamiliar letter string can elicit.

Although previous studies have shown that pseudowords can be indicative of meaning, whether the semantic underpinnings of words and pseudowords rely on the same cognitive processes, and hence whether they both activate the same semantic system, is currently not known. Indeed, previous studies mainly asked participants to make overt or covert decisions on whether a target string of letters is a correctly spelled word in a reference language, as in the lexical decision tasks (e.g., Hendrix & Sun, 2021). Yet, in a simple lexical decision paradigm, pseudowords are presented in isolation and it is thus not possible to probe whether their elaboration would tap into the same source of semantic information subserving word processing; rather, semantic effects for pseudowords as the ones

observed by Hendrix and Sun (2021) could emerge from ad-hoc, strategic processes triggered by the attempt of making sense of the meaningless, with not much in common with the comprehension of actual words. In order to put this aspect to test, here we consider whether the semantic processing of a word can modulate a subsequent pseudoword-elicited meaning activation. We relied on the priming task, in which participants are shown two strings of letters presented sequentially one after the other (with the former being referenced as the *prime* and the latter as the *target*) and are then asked to indicate whether the latter is a word or a pseudoword. Priming tasks can be used to investigate the effects induced by the semantic similarity between the prime and target words (Draine & Greenwald, 1998; Xavier Alario et al., 2000). For instance, it has been shown that when prime and target word pairs are semantically related (e.g., *doctor – physician*), participants’ responses are faster compared with when they are unrelated (e.g., *doctor – window*). Crucially, DSMs have been shown to be particularly efficient in predicting participants’ performance in priming tasks (Günther et al., 2016; Jones et al., 2006; Marelli, 2017). In DSMs, each word is represented by a vector indexing a position within a common semantic space: hence, the closer the semantic distance between the vector activated by the prime word and the one activated by the target word, the faster the participants’ response.

In priming tasks, researchers have mostly investigated semantic effects by looking at participants’ performance when judging (meaningful) verbal stimuli (i.e., words). While previous studies have demonstrated that both words and pseudowords can prime orthographically similar items (Rueckl & Olds, 1993; Stark & McClelland, 2000), the idea of a semantic priming induced on target pseudowords by prime words was simply inconceivable from both theoretical (i.e., pseudowords do not have any meaning and, therefore, do not have corresponding representations in semantic memory) and practical (i.e., it was impossible to operationalize whether a pseudoword is semantically related or not to a word) standpoints. The priming task is thus the ideal tool to investigate whether the processing of words and pseudowords can be traced back to common semantic processes and whether participants “search” for a meaning for both words and pseudowords within a common semantic space.

In Experiment 1, we first analyzed the data of the Semantic Priming Project (Hutchison et al., 2013), the largest database available of priming data, by predicting participants’ response latencies and accuracy using semantic relatedness values between the word prime and the pseudoword target as extracted from *fastText*. In Experiment 2, we ran an independent replication on a new language (Italian), critically controlling for a series of confounding variables (e.g., orthographic similarity and morphological structure). If the processing of words and pseudowords is supported by the same associative mechanisms operating over a common semantic space, we should expect a more impaired

performance for pseudowords that are more semantically related to prime words, as the larger semantic activation would make it more difficult to judge the target as a pseudoword. That is, when participants have to judge the meaning of a pseudoword, their performance should be affected by the semantic pre-activation elicited by the prime word, with such interference being more pronounced when the pattern activated by the prime word overlaps more with the one activated by the target pseudoword.

Experiment 1: Analysis of the Semantic Priming Project megastudy

Methods

Materials

We analyzed lexical-decision data from the semantic priming project (SPP; Hutchison et al., 2013), a megastudy including response latencies from 768 subjects for 1,661 English target words and 1,665 target pseudowords, preceded by different prime words. In the SPP, each experimental trial began with a fixation cross presented in the center of the screen, followed by a prime word in uppercase letters. Next, a blank screen appeared, creating a 200- or 1,200-ms stimulus onset asynchrony (SOA). Finally, a lowercase target string was presented (i.e., it could either be a word or a pseudoword). Participants had to indicate whether the target string was a word or a pseudoword.

In the SPP, pseudowords were generated by changing one or two letters from each target word to form a pronounceable pseudoword. From these data we excluded the prime-target pairs in which the prime or the target had three or fewer letters and those for which a limited number of datapoints were available (i.e., 10 or less; $< .0001\%$ of the trials was excluded). This latter choice was driven by the need of a relatively large numerosity in order to compute reliable mean accuracy and reaction times. Following the application of this criterion, the number of stimuli eventually included as targets in our analysis was 2,987 (1,491 words and 1,496 pseudowords), with the number of unique prime-target pairs being 16,732 in the short-SOA condition and 16,543 in the long-SOA condition.

Distributional semantic model

The DSM used here was *fastText* (Joulin et al., 2016), and words and pseudowords vectors were retrieved from the English 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, 1993; and realized computationally by Bojanowski

et al., 2017) to take into account sub-word information by inducing semantic representations as the sum of the vectors of the letter n-grams associated with each word. That is, *fastText* computes the semantic representation of each string of letters as the sum of the vector of the full string plus all the vectors of the 3- to 6-grams that compose it. As an example, consider the word <asking>, composed by different-length character n-grams, as reported in the Table 1. The *fastText*-induced representation will be the sum of the <asking> word vector along with the vectors of all the elements reported in Table 1.

A similar approach can be applied to unattested strings in order to capture the semantic information associated to 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 on 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 mapped 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 (around 630 billion words) and Wikipedia (around 9 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 $w_{-n}, \dots, w_{-1}, w_1, \dots, w_n$.

Using *fastText*, we therefore obtained semantic representations for both the words and the pseudowords included in this analysis (i.e., the prime-target pairs). For each pair we computed a semantic-relatedness index (SRel) based on the cosine of the angle formed by vectors representing the meanings of the corresponding strings. The higher the SRel value, the more semantically related the letter strings are expected to be, as estimated by the model.

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 used in the SPP. 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. We found that a full vector representation was available for 390 of the considered pseudowords. To overcome possible biases induced by the

availability of this further piece of information, we removed such whole-pseudoword vectors from the .bin file used to retrieve the n-gram representations. For these 390 pseudowords, vectors were computed on the basis of this updated semantic space.

Data analysis

The codes for reproducing the analyses are made fully available at: <https://osf.io/pt97n/>; while the SPP data are available at: <https://www.montana.edu/attmemlab/spp.html>. The statistical analyses were performed using *R-Studio* (RStudio Team, 2015). The plots reported were obtained using the *effects* R package (Fox, 2003; Fox & Weisberg, 2019). The analyses were performed at the item-level by computing the mean correct reaction times (RTs) and the mean accuracy for each prime-target pair, separately for the two SOAs, in line with previous studies analyzing large databases of this kind (e.g., Amenta et al., 2020a). Trials in which participants' RTs were faster than 300 ms or slower than 1500 ms were removed from the computation of the mean RTs (3% of the trials excluded).

For each prime-target pair we further computed the corresponding Levenshtein distance (hence OrthoDist), which measures the orthographic distance between two strings of symbols by quantifying the minimum number of single-character edits (e.g., insertions, deletions, or substitutions) required to change one element into the other (i.e., the prime into the target). The OrthoDist was computed using the *stringdist* R package (Van der Loo, 2014).

For each prime and for each target, we computed their length (i.e., the number of letters in the letter string) and, as a measure of orthographic neighborhood (hence OrthoNeigh), we computed its OLD20 (i.e., the average OrthoDist between the letter string and its 20 closer neighbors; Yarkoni et al., 2008). In particular, the OrthoNeigh was computed using the *vwr* R package (Keuleers, 2013) and, to select the possible neighbors, we considered the 15,000¹ most frequent words in SUBTLEX-us (Brysbaert & New, 2009). The inclusion of these covariates ensures that our effects are not trivially determined by the orthographic overlap between prime and target. In fact, since *fastText* builds primarily on orthographic n-grams, it might be the case that it implicitly captures the surface similarity between words, over and above their association in terms of meaning. Evaluating the effect of OrthoDist and OrthoNeigh allows to control for this possibility.

¹ The decision to include only the 15,000 most frequent words was made i) considering the need to limit the overestimation of the (possible) effect of low frequency words in the computation of the predictor and ii) following Hendrix & Sun (2021), who limited the overall words used to compute the OLD20 to the 18,000 stimuli included in the dataset they were analyzing.

Following Hendrix & Sun (2021), for each target stimulus we also computed an index labeled semantic neighborhood (hence SNeigh). SNeigh was computed as the mean semantic similarity between each pseudowords and its 5 closest neighbors within the 15,000 most frequent words in SUBTLEX-US. Higher values thus indicate that a certain pseudoword is more embedded in the semantic space. Similarly, following Hendrix and Sun (2021), for each target stimulus we also computed an index labeled mean binomial frequency (hence MBF). MBF was computed as the average frequency of the letter bigrams that compose each verbal stimulus. Letter bigram frequencies were retrieved from Google Corpus Data².

Finally, for each target we included a predictor labeled Target-related frequency. Such variable was differently operationalized depending on the considered type of target stimulus (word vs. pseudoword). Specifically, for word targets this predictor consists in the log-frequency of the word as retrieved from SUBTLEX-US, while for pseudoword targets this predictor refers to the log-frequency of its first orthographic neighbor (in terms on Levenshtein distance) within the 15,000 most frequent words in the SUBTLEX-US.

The predictors included in our statistical models were thus: SRel, OrthoDist, Prime length, Target length, Prime OrthoNeigh, Target OrthoNeigh, Prime frequency retrieved from the SUBTLEX-us (Brysbaert & New, 2009), MBF, SNeigh, Target-related frequency and the type of target stimulus (word vs. pseudoword).

To investigate the effect of SRel we estimated four linear models (LMs), separately for short and long SOA and for RTs or accuracy as dependent variables (DVs). To obtain gaussian data distributions, RTs were log-transformed, while for accuracy we applied the arcsine-square-root transformation (Kruskal, 1968). The decision to log-transform RTs was driven by the fact that their distribution was positively skewed, while accuracy was negatively skewed.

In the LMs estimated we included SRel, OrthoDist, Prime length, Target length, Prime OrthoNeigh, Target OrthoNeigh, Prime frequency, MBF, SNeigh, and Target-related frequency as continuous predictors, and type of target (word vs. pseudoword) as categorical predictor along with its interactions with all the other variables. Additionally, we also included the three-way interaction SRel by Target-related frequency by type of target (along with the relative two-way interaction SRel by Target-related frequency). Such interaction was aiming at controlling the possible impact of pseudoword neighbors: pseudowords included in the experiment are in fact highly similar to existing English words (since in the SPP pseudowords were created by changing one or two letters from existing English words). If the SRel effect is dependent on the similarity of the pseudowords with such very similar, existing

² See: <https://gist.github.com/lydell/c439049abac2c9226e53>

English words, the three-way interaction should be significant in the final model. Conversely, if the three-way interaction is not significant, the eventual SRel by type of target effect would be dependent on the specific semantic characteristics evoked by the pseudowords themselves (i.e., the distributional pattern of their sublexical elements in natural language). In particular, the models tested was:

$$\begin{aligned} DV \sim & (SRel + OrthoDist + Prime\ length + Target\ length + Prime\ OrthoNeigh \\ & + Target\ OrthoNeigh + Prime\ frequency + MBF + SNeigh \\ & + Target\text{-}related\ frequency) \times Type + SRel:Target\text{-}related\ frequency \\ & + SRel:Target\text{-}related\ frequency:Type \end{aligned}$$

We then performed a model selection using the *MuMIn* R package, with the function *dredge* (Bartoń, 2020). This procedure selects the best fitting model (i.e., the one with the lowest Akaike information criterion, which returns an estimation of the quality of the model, AIC; Akaike, 1973) fitting all the possible combinations of the fixed effects included. In order to check for the consistency of the effects, we also estimated all the models with $\Delta AIC < 2^3$, as models in this AIC range can be considered as equivalent in explaining participants' performance (Hilbe, 2011). Here we report only results related to the interaction of interest between SRel and type of target; for the full results (including the effects of the other predictors) please see the Supplementary Materials.

Results

Descriptive statistics

The correlation matrix between the predictors tested for both short and long SOAs is reported in Figure 1 (correlational values were computed at the item level). Overall, the correlations ranged from very low, as the one between SRel and OrthoDist ($r < .01$), to high, as the ones between target length and target OrthoNeigh ($r = .90$).

Reaction times

Short SOA

³ All the *dredge* outputs and the results of the models with $\Delta AIC < 2$ are reported at: <https://osf.io/pt97n/>.

The best model identified by the *dredge* selection was the full model, $R^2 = .35$; $Adj. R^2 = .35$ (Figure 2A). Three models had $\Delta AIC < 2$ compared with the best model selected and reported comparable interactions and marginal means of linear trends effects.

In particular, in the best model, the three-way interaction SRel by Target-related frequency by type of target was not significant, $F(1, 16708) = 2.32, p = .12$, while the interaction between SRel and type of target was significant, $F(1, 16708) = 56.27, p < .001$, indicating that, for word targets, the higher the SRel with the word prime, the lower the time needed to correctly recognize the stimulus, $t(16708) = -13.20, p < .001, b = -.08$. Conversely, for pseudoword targets, the higher the SRel with the word prime, the higher the time needed to correctly reject the stimulus, $t(16708) = 2.36, p = .01, b = .04$. Notably, the same effect for pseudowords targets was present also in the three other models with $\Delta AIC < 2$.

Long SOA

The best model identified by the *dredge* selection included the effects of SRel, Target length, Target OrthoNeigh, Prime frequency, MBF, Sneigh, Target-related frequency and their interactions with the type of target, and, additionally, the interaction Srel by Target-related frequency, $R^2 = .30$; $Adj. R^2 = .30$ (Figure 2B). Eight models had $\Delta AIC < 2$ compared with the best model selected and reported comparable interactions and marginal means of linear trends effects.

In particular, in the best model, the interaction between SRel and type of target was significant, $F(1, 16526) = 22.21, p < .001$, indicating that, for word targets, the higher the SRel with the word prime, the lower the time needed to correctly recognize the stimulus, $t(16526) = -10.54, p < .001, b = -.06$. Conversely, for the pseudoword targets, the effect of SRel was not significant, $t(16526) = .28, p = .78, b = .005$ (nor it was in the other eight models tested with $\Delta AIC < 2$).

Accuracy

Short SOA

The best model identified by the dredge selection included all the predictors and interactions tested with the exception of the three-way interaction SRel by Target-related frequency by type of target, $R^2 = .06$; $Adj. R^2 = .06$ (Figure 2C). Nineteen models had $\Delta AIC < 2$ compared with the best model selected and reported comparable interactions and marginal means of linear trends effects.

In particular, in the best model, the interaction between SRel and type of target was significant, $F(1, 16709) = 38.35, p < .001$, indicating that, for word targets, the higher the SRel with the word prime, the higher the accuracy, $t(16709) = 5.26, p < .001, b = .06$. Conversely, for the pseudoword targets, the higher the SRel with the word prime, the lower the accuracy, $t(16709) = -2.86, p = .004, b = -.10$. Notably, the same effect for pseudoword targets was present also in the nineteen other models with $\Delta AIC < 2$.

Long SOA

The best model identified by the *dredge* selection included the effects of SRel, Target length, Target OrthoNeigh, Prime frequency, MBF, SNeigh, Target-related frequency and their interactions with the type of target, $R^2 = .05$; *Adj. R*² = .05 (Figure 2D). Five models had $\Delta AIC < 2$ compared with the best model selected and reported comparable interactions effects, while for marginal means of linear trends effects, they were consistent across target words, but not fully consistent across target pseudowords.

In particular, in the best model, the interaction between SRel and type of target was significant, $F(1, 16527) = 32.80, p < .001$, indicating that, for word targets, the higher the SRel with the word prime, the higher the accuracy, $t(16527) = 7.53, p < .001, b = .07$. Conversely, for the pseudoword targets, the higher the SRel with the word prime, the lower the accuracy, $t(16527) = -2.05, p = .04, b = -.07$. Note that this latter effect for pseudoword targets was not significant in two of the other five models with $\Delta AIC < 2$.

Discussion of Experiment 1

In Experiment 1 we analyzed data from the SPP (Hutchinson et al., 2013) and found results consistent with the idea that the processing of words and pseudowords can influence each-other at the semantic level. In particular, at a short SOA (i.e., 200 ms) the time needed to judge whether a target was a pseudoword increased as a function of the semantic relatedness with the prime word, while the accuracy decreased, indicating in turn that the decision was more difficult the more the prime word was semantically related to the pseudoword target. Hence, the meaning activated by the prime word interfered with the semantics activated by the pseudoword, suggesting that this interference occurred via shared semantic information. At a long SOA (i.e., 1200 ms) no SRel effect was found for pseudowords on RTs, while for accuracy results were not fully consistent across all the models estimated. These results may indicate that the semantic effect for pseudowords may be the result of

automatic processes operating at shorter SOAs, while the effect may disappear at longer SOAs when controlled processing strategies become activated (e.g., Küper & Heil, 2009). This is consistent with previous studies showing that weaker effects in semantic priming might be associated with faster activation decay, and hence are less likely to be observed in long-SOA paradigms (e.g., Petilli et al., 2021).

Yet, the pseudowords included in the SPP were not selected ad hoc for our aims and may contain key linguistic units trivially known to carry meaning. Indeed, the morphological characteristics of the strings of letters composing a pseudoword, such as particular linguistic regularities like suffixes (e.g., in English words ending in *-acy* are indicative of a state or a quality, such as in the words *privacy*, *fallacy* or *delicacy*), are informative of meaning and were shown to affect processing (e.g., Amenta, et al., 2020b; Marelli & Baroni, 2015; Ulicheva et al., 2021). As a consequence, the observed semantic effect could have been driven by these linguistically defined meaning-carrying sublexical elements, and not by a broader, more general, sub-word semantic effect. Additionally, SPP pseudowords were obtained by changing one or two letters of a real English word (e.g., *abnolmal*, *alnohol*, *belax*) and were hence highly similar to existing words; the observed semantic effect could have therefore been driven by some sort of turple effect (e.g., Forster & Hector, 2002) captured by the structure of *fastText*. Even if this is unlikely considering the lack of a significant three-way interaction between SRel, target type and Target-related frequency, a tighter control of these potential influences would be desirable. On these grounds, we opted to perform an independent replication in a different language (i.e., Italian), by taking into account these potential confoundings. In particular, in Experiment 2, we built a new set of pseudowords excluding items i) that included suffixes and ii) selecting a priori word-pseudoword pairs across a wider semantic-relatedness distribution.

Experiment 2

Methods

Participants

Seventy-five students participated in this experiment for course credits (22 males, M age = 23.36 years, $SD = 2.91$). We recruited as many students as possible before the end of the semester. 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

As in the analysis of the data from the SPP, we employed *fastText* (Joulin et al., 2016). The only difference is that in this case we considered the Italian word vectors (Grave, et al., 2018). The model was trained on Common Crawl and Italian Wikipedia (around 11 billion words) using CBOW with 300 dimensions, character n-grams of up to length 5, and a window of size 5.

As for Experiment 1, we checked whether a full vector was available in the corpus for the pseudowords used. We found that for 12 pseudowords such a representation was available. Thus, to overcome possible effects related to these whole-pseudoword vector, we removed them from the .bin file used to compute the pseudoword representations. For these 12 pseudowords, the vectors were thus based on this updated semantic space.

Stimuli

Stimuli were prepared starting from the words included in the Italian ANEW database (Montefinese et al., 2014). In the study by Montefinese and colleagues, a large number of participants were asked to rate Italian words on different scales, including concreteness (ranging from 1, abstract, to 9, concrete).

First, from the ANEW database, we selected the 400 most concrete words: this was aimed at reducing the variability between the included stimuli, since concrete words are generally understood more easily compared with abstract ones (e.g., Kounios & Holcomb, 1994; for a similar approach, in which concrete words were used in a semantic priming task, see: Günther et al., 2016). Then, using Wuggy (Keuleers, & Brysbaert, 2010), we automatically obtained 3570 pseudowords. Wuggy was set using its standard parameters, that is: orthographic Italian module, restricted match length of subsyllabic segments, restricted match letter length, restricted match transition frequencies, match segments 2 out of 3. Starting from a given word, Wuggy allows for the generation of written polysyllabic pseudowords that obey a given language's phonotactic constraints and that match its template in subsyllabic structure. That is, Wuggy generates highly word-like pseudowords but also stimuli that are not easily identifiable as related to existing words.

Subsequently, using *fastText*, we induced vector representations for each of the 400 words selected from the ANEW database and for the obtained pseudowords. We thus paired each word with each pseudoword, and words with each other, obtaining 1,428,000 word-pseudoword pairs and 159,600 word-word pairs, and computed a semantic-relatedness index (SRel) for each pair as in Experiment

1. Then, word-pseudoword and word-word pairs were included as prime-target stimuli based on the following criteria: i) pseudowords had to be orthotactically legal for Italian; ii) pseudowords had not to include recognizable suffixes⁴ (e.g., the words that ended with -aio, -iato, -ista, etc.⁵ were not included); iii) to allow for an uniform-like distribution of SRel values (applied to both word-word and word-pseudoword pairs; with SRel ranging from .007 to .49) we included a similar amount of pairs for each .1 step of SRel (see the Supplementary Materials for the graphical distributions). Additionally, each letter string (being either a word or a pseudoword) had to appear maximum one time in the whole set of stimuli. These criteria guarantee that potential effects are not driven by semantic information carried by morphological cues (Marelli & Baroni, 2015), and depend only on the semantic impact of sub-lexical, non linguistically-defined, chunks.

The final set of stimuli included 180 pairs (90 word-word pairs and 90 word-pseudowords pairs).

Procedure

Participants were shown two letter strings, presented sequentially one after the other, and were asked to judge whether the latter stimulus was an Italian word or not. The former word was presented in uppercase letters (see: Hutchison et al., 2013) and participants were required to respond only to the latter (presented in lowercase letters), as fast and as accurately as possible by pressing the A or L key using the left or right index fingers, respectively; the response keys were counterbalanced across participants. The trials were presented in random order.

Hence, Experiment 2 was identical with the task employed in the SPP except for two methodological details. Firstly, here participants were tested using PsychoPy (Peirce et al., 2019) through the online platform Pavlovia (<https://pavlovia.org/>). Secondly, here we employed a fixed SOA. Specifically, each trial started with a central fixation cross (presented for 500 ms) and was followed by a first word (the prime, presented for 500 ms), by a brief interval (100 ms) and then by a second stimulus (the target word or pseudoword, presented for 3000 ms). Participants' response ended the trial and moved to a blank screen (presented for 1000 ms), which preceded the fixation cross of the next trial.

⁴ 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 (e.g., *anti-*, *contro-*, *extra-*) was included in our set of stimuli.

⁵ For a complete list of Italian suffixes see: [https://www.treccani.it/enciclopedia/suffissi_\(La-grammatica-italiana\)/](https://www.treccani.it/enciclopedia/suffissi_(La-grammatica-italiana)/)

Data analysis

The statistical analyses were performed using *R-Studio* (RStudio Team, 2015) and the *lme4* R package (Bates et al., 2015). Plots reported were obtained using the *effects* R package (Fox, 2003; Fox & Weisberg, 2019).

The dependent variables were accuracy (i.e., “incorrect” responses were scored as 0s, whereas “correct” responses as 1s) and log-transformed correct reaction times (RTs). Trials in which participants’ RTs were faster than 300 ms or slower than 1,500 ms were removed from the analysis (3% of the trials excluded).

As in Experiment 1, for each prime-target pair we computed the OrthoDist using the *stringdist* R package (Van der Loo, 2014). For each prime and target, we also computed its length, OrthoNeigh with the *vwr* R package (Keuleers, 2013), and, for each target, SNeigh. In this case, we used the 15,000 most frequent words in the Italian SUBTLEX (<http://crr.ugent.be/subtlex-it/>). Similarly, MBF was computed for each target. Letter bigram frequencies were retrieved from ItWaC Data⁶. Finally, as for Experiment 1, for each target we included the Target-related frequency predictor, this time computed using SUBTLEX-IT.

The predictors in our statistical models were identical to those included in Experiment 1. To investigate if the SRel between the prime word and the target could affect RTs, we ran a Linear Mixed Model (LMM) having RTs as dependent variable, while to probe whether there was an effect on accuracy, we ran a generalized LMM (GLMM).

In the GLMM we set the number of adaptive Gauss–Hermite quadrature points to zero (i.e., setting nAGQ = 0, see: Bates et al. 2015) for computational efficiency reasons⁷. In both models, participants and prime-target pairs were included as random intercepts. In both models, the fixed effects estimated were identical to Experiment 1. In particular, in the *lme4* syntax, the models tested were:

$$\begin{aligned}
 DV \sim & (SRel + OrthoDist + Prime\ length + Target\ length + Prime\ OrthoNeigh \\
 & + Target\ OrthoNeigh + Prime\ frequency + MBF + SNeigh \\
 & + Target\text{-}related\ frequency) \times Type + SRel:Target\text{-}related\ frequency \\
 & + SRel:Target\text{-}related\ frequency:Type + (1|Participant) + (1|Item)
 \end{aligned}$$

⁶ See: https://corpora.dipintra.it/public/run.cgi/first_form?corpname=itwac_full;align=

⁷ This procedure allows for a relatively fast estimation of the models, but it is less optimized compared with other methods. We opted for this procedure as other optimizers would have been extremely time-consuming. Additionally, control analyses testing the best models with both nAGQ = 0 and nAGQ = 1 showed that the $\Delta AICs$ were overall similar.

For both models we performed a model selection using the *MuMIn* R package (Bartoń, 2020), in analogy with the analyses on the SPP data. In order to check for the consistency of the effects, we also estimated all the models with $\Delta AIC < 2^8$.

Finally, to exclude the impact of overly influential outliers, after having fitted the best model identified by the *dredge* selection on RTs, data points were removed on the basis of a threshold of 2.5 *SD* standardized residual errors (model criticism; Baayen, 2008). Results based on the refitted models are reported. Here we report only results related to the interaction SRel by type of target, for the full results (including the effects of the other predictors) please see the Supplementary Materials.

Results

Descriptive statistics

The correlation matrix between the predictors is reported in Figure 3. Overall, the correlations ranged from very low, as the one between SRel and MBF ($r = .03$), to high, as the one between prime length and prime OrthoNeigh ($r = .90$).

Reaction times

The best model identified by the dredge selection included SRel, Target-related frequency and their interaction with type of target and, additively, Prime OrthoNeigh, $Pseudo-R^2$ (total) = .49; $Pseudo-R^2$ (fixed effects) = .13. No other model had $\Delta AIC < 2$ compared with the best model selected.

The interaction between SRel and type of target was significant, $F(1,170) = 16.98, p < .001$, indicating that the effect of SRel differently predicted participants' RTs across the two types of targets (Figure 4A). That is, for the word targets, the higher the SRel, the faster the RTs, $z = -3.24, p < .001, b = -.12$; while for the pseudoword targets, the higher the SRel, the slower the RTs, $z = 2.61, p = .009, b = .10$.

⁸ All the *dredge* outputs and the results of the models with $\Delta AIC < 2$ are reported at: <https://osf.io/pt97n/>.

Accuracy

The best model included SRel, Target length, Target OrthoNeigh, Target-related frequency and their interaction with the type of target and, additively, Prime OrthoNeigh, $Pseudo-R^2$ (total) = .34; $Pseudo-R^2$ (fixed effects) = .08. Twenty-one models had $\Delta AIC < 2$ compared with the best model selected. The interaction SRel by type of pair was included in sixteen over twenty-two models and it was significant in twelve (with the effect of SRel on pseudowords being significant in eleven of them).

In particular, in the best model, the interaction SRel by type of pair was significant, $\chi^2(1) = 4.34$, $p = .03$, indicating that the effect of SRel differentially predicted participants' accuracy across the two types of targets. For pseudowords the higher the SRel, the lower the accuracy, $z = -2.08$, $p = .03$, $b = -1.33$; while for words, no significant effects were found, $z = 1.17$, $p = .24$, $b = .81$ (Figure 4B), with this effect being likely caused by ceiling effects (i.e., participants' accuracy for words was overall very high).

Discussion of Experiment 2

In Experiment 2 we replicated the semantic-priming effect on pseudowords (and, specifically, the effect observed at the shortest SOA in Experiment 1) for both reaction times and accuracy. In particular, reaction times increased for increasing semantic relatedness between word primes and pseudoword targets. Consistent effects were observed here also for accuracy, with increasing semantic relatedness associated with a higher number of errors (i.e., pseudoword targets were incorrectly identified as words), despite in this case the pattern was less consistent as it was not present in all the models estimated. This testifies a pervasive influence of the semantic association between prime and target stimuli on the overall performance, regardless of whether these stimuli were part of the familiar lexicon or not.

Notably, these effects were observed although we controlled for various confounding variables (e.g., excluding pseudowords with Italian suffixes). As a consequence, the observed semantic effects are to be understood in terms of a general form-meaning systematicity occurring at the sub-word level. That is, the presentation of a prime word activates a pattern of meaning in participants' semantic memory; when this pattern overlaps with the one elicited by the sublexical information of the target pseudoword, participants' performance is affected, increasing the response for this latter.

General Discussion

In the present study, we probed for the first time whether the same associative mechanisms governing word meaning also subserve the processing of pseudowords. We took advantage of a distributional semantic model, namely *fastText* (Bojanowski et al., 2017), able to approximate the semantic information carried by sub-word units in language and explored whether this information can predict participants' performance in a typical priming task. In Experiment 1, we first analyzed the data of the semantic priming project (Hutchison et al., 2013), while in Experiment 2 we ran an independent replication on a different language, controlling for a series of potential confounding variables. Results were consistent across the two experiments and showed that at increasing semantic relatedness between prime word and target pseudoword (i.e., the more the corresponding orthographic elements elicited an overlapping semantic pattern) participants' reaction times increased, while accuracy decreased (although the pattern for accuracy was less consistent). Notably, as we used a lexical decision task, participants were not explicitly asked to link the information carried by pseudowords with aspects of meaning. That is, we could directly probe whether the (semantic) activation elicited by the prime word would *automatically* interfere with pseudoword processing, with no need to encourage a semantic interpretation of the unfamiliar string. This suggests that, when encountering any (well-formed) linguistic string, a speaker will routinely activate information in semantic memory, irrespective of the lexicality of the stimulus. Such information would tap into the same system storing the meanings of actual words, leading to the possibility of semantic priming effects between word and pseudowords. This speaks for a natural inclination, in human beings, to make sense of their experience, even when this latter is, at face value, meaningless: when the whole string does not have an explicit link with a meaning stored in semantic memory, humans can still rely on the distribution of sub-word elements in linguistic contexts to derive a semantic intuition. This interpretation can be seen as the linguistic counterpart of well-known visual illusory effects such as pareidolia and apophenia (i.e., the tendency of identifying meaningful patterns where none truly exist, as when people tend to see faces from objects or meaningless visual patterns), with all these phenomena possibly tapping on a common substrate of human's tendency to search for a meaning in complex and fuzzy environments.

These findings well complement non-arbitrary perspectives on language processing (Dingemanse et al., 2015). The semantic effect found for pseudowords can indeed be explained with reference to the humans' tendency to detect systematic and statistical regularities in the (language) environment (Romberg & Saffran, 2010; Vidal et al., 2021), here expressed in the form of nuanced distributional patterns of sub-word information. Accordingly, seminal studies have shown that children can extract syllabic patterns based on transitional probabilities between elements, when these patterns are

presented in a continuous stream (Saffran et al., 1996). Subsequent studies have demonstrated that humans extensively exploit these experienced-based learning mechanisms in grammatical, phonological, and orthographical processing (e.g., Adriaans, & Kager, 2010; Altmann, et al., 1995; Saffran & Wilson, 2003; Thompson & Newport, 2007; and for a review: Bogaerts et al., 2021; Christiansen, 2019), and similar effects have been reported also at for semantic processing (e.g., Ulicheva et al., 2020). Integrating this evidence, our findings indicate that humans could exploit form-to-meaning regularities at a deeper level (i.e., in term of distributional patterns of sub-word information), with these processes possibly being far more extensive than previously thought (for recent evidence see: Chuang et al., 2021). Specifically, these findings support seminal connectionist theories (e.g., Harm & Seidenberg, 1999; Plaut & Gonnerman, 2000; Seidenberg & Gonnerman, 2000) as well as more recent perspectives (Baayen et al., 2011), which map form onto meaning relying on the parallel nature of neural processing, and the distributed nature of neural representations. Accordingly, units in language would primarily depend on the way they are acquired, organized and used by humans (Marelli et al., 2020).

Additionally, the present findings on pseudowords processing constitute a serious challenge to classical theories on priming tasks. Seminal accounts explained participants' performance in the semantic priming task by maintaining that the processing of the first word automatically pre-activates its semantic neighbors (i.e., the semantic activation of the prime spreads through its neighborhood in a network in which each word is a unit): hence, the semantically closer the prime and the target words in the semantic network, the faster the judgement about the target (Collins & Loftus, 1975). Following these accounts, a performance impairment when judging pseudowords induced by their semantic relationship with a prime word would be simply inconceivable, since pseudowords would not have any place in such symbolic semantic network. Contrarily to these views, our results show that such semantic effects are possible, suggesting that distributional models of semantic memory could be more well-fitted to explain the structure of human semantic memory as compared to network-based models (Günther et al., 2019).

Interestingly, several studies have shown facilitatory effects of iconicity and systematicity on word acquisition (e.g., Cassani & Limacher, 2021; Imai et al., 2015; Monaghan et al., 2012). These pieces of evidence suggest that, in vocabulary expansion, the ease of attaching a meaning to the newly acquired label plays a substantial role, and hence that the semantic system provides an important contribution to such process. The present results expand and further support these conclusions, by showing that, even in a context in which there is no new word to be acquired, processing unfamiliar strings implies the activation of semantic information that is entangled with familiar words. In fact, pseudowords in our lexical decision task are not characterized as meaning-bearing units that are

(potentially) part of the lexicon, as new words during acquisition are assumed to be; on the contrary, they are defined in the very experiment instructions as non-existing terms, that must be identified as such by the participants. Nevertheless, they are shown to be affected by the very semantic information activated by the lexical prime, speaking for how pervasive and automatic the meaning-extraction process is, even in front of a string of symbols that has been clearly framed as not being part of the lexicon of a language, and providing new evidence in favor of the role of systematicity in word processing.

Given that this evidence was crucially obtained by relying on model building on linguistic distributions, the present results are in line with the idea that human semantic memory has developed by exploiting and re-cycling general-purpose associative learning mechanisms (Günther et al., 2019), with increasing evidence showing that reward prediction error principles are at play in the formation of declarative memory traces (Ergo et al., 2020). Specifically, these processes seem to apply to low-level, featural elements in the environment, in the present case by keeping track and ultimately quantifying the learning history of sub-word units in language. The distributional pattern of sub-word information would be then combined when elaborating more complex elements (i.e., words or pseudowords), with such combination representing a meaningful cue built on a linguistic experiential basis. The evidence presented here provides direct support to such a distributed model of human semantic memory, in that each word would activate a representation based on the distributional pattern of its sub-units. Notably, these domain-general processes would be employed for any linguistic stimulus and, consequently, also for new words, which would end up being represented by the very same memory system. The involvement of these flexible cognitive processes is ideal considering that semantic memory (and language more generally) is a generative system, that must deal with novel information almost on a daily basis. Indeed, an adult speaker learn from 1.7 (Nation, 2006) to 11 (Nagy & Anderson, 1984) new words per day (for a complete discussion see: Brysbaert and colleagues, 2016).

In conclusion, using distributional semantic models we provide evidence for a semantic priming effect on pseudowords, demonstrating that semantic induction from words vis-à-vis pseudowords can be traced back to domain-general associative mechanisms operating over the same memory system. Our findings directly support theories on the non-arbitrariness of language and provide novel insights on 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.
- Akaike, H. (1973). Information theory and an extension of the maximum likelihood principle. In B. N. Petrov & F. Caski (Eds.), *Proceedings of the Second International Symposium on Information Theory* (pp. 267-281). Budapest: Akademiai Kiado.
- 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.
- Amenta, S., Crepaldi, D., & Marelli, M. (2020a). Consistency measures individuate dissociating semantic modulations in priming paradigms: A new look on semantics in the processing of (complex) words. *Quarterly Journal of Experimental Psychology*, 73(10), 1546-1563.
- Amenta, S., Günther, F., & Marelli, M. (2020b). A (distributional) semantic perspective on the processing of morphologically complex words. *The Mental Lexicon*, 15(1), 62-78.
- 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.
- Baayen, R.H. (2008). *Analyzing linguistic data: A practical introduction to statistics using R*. Cambridge: Cambridge University Press.
- Bartoń, K. (2020). MuMIn: Multi-Model Inference. R package version 1.43.17. <https://CRAN.R-project.org/package=MuMIn>
- Bates, D., Maechler, 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.
- 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.
- Brysbaert, M., & New, B. (2009). Moving beyond Kučera and Francis: A critical evaluation of current word frequency norms and the introduction of a new and improved word frequency measure for American English. *Behavior Research Methods*, 41(4), 977-990.
- 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, 1116.
- Cassani, G., & Limacher, N. (2021). Not just form, not just meaning: Words with consistent form-meaning mappings are learned earlier. *Quarterly Journal of Experimental Psychology*, 17470218211053472

- 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(3), 945-976.
- De Saussure, F. (1916: 2011). *Course in General Linguistics*. Columbia University Press.
- 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.
- Draine, S.C., & Greenwald, A.G. (1998). Replicable unconscious semantic priming. *Journal of Experimental Psychology: General*, 127(3), 286.
- Ergo, K., De Loof, E., & Verguts, T. (2020). Reward prediction error and declarative memory. *Trends in Cognitive Sciences*, 24(5), 388-397.
- Forster, K. I., & Hector, J. (2002). Cascaded versus noncascaded models of lexical and semantic processing: The turtle effect. *Memory & Cognition*, 30(7), 1106-1117.
- Fox, J. (2003). Effect Displays in R for Generalised Linear Models. *Journal of Statistical Software*, 8(15), 1-27. URL: <http://www.jstatsoft.org/v08/i15/>
- Fox, J., & Weisberg, S. (2019). *An R Companion to Applied Regression*, 3rd Edition. Thousand Oaks, CA <http://tinyurl.com/carbook>
- Gatti, D., Rinaldi, L., Marelli, M., Mazzoni, G., & Vecchi, T. (2021). Decomposing the semantic processes underpinning veridical and false memories. *Journal of Experimental Psychology: General*.
- Grave, E., Bojanowski, P., Gupta, P., Joulin, A., & Mikolov, T. (2018). Learning word vectors for 157 languages. arXiv preprint arXiv:1802.06893.
- Günther, F., Dudschig, C., & Kaup, B. (2016). Predicting lexical priming effects from distributional semantic similarities: A replication with extension. *Frontiers in Psychology*, 7, 1646.
- 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. (1999). Phonology, reading acquisition, and dyslexia: insights from connectionist models. *Psychological Review*, 106(3), 491.
- 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.

- Hilbe, J. M. (2011). *Negative binomial regression*. Cambridge University Press.
- Hutchison, K. A., Balota, D. A., Neely, J. H., Cortese, M. J., Cohen-Shikora, E. R., Tse, C. S., ... & Buchanan, E. (2013). The semantic priming project. *Behavior Research Methods*, 45(4), 1099-1114.
- Imai, M., Kita, S., Nagumo, M., & Okada, H. (2008). Sound symbolism facilitates early verb learning. *Cognition*, 109(1), 54-65.
- Imai, M., Miyazaki, M., Yeung, H. H., Hidaka, S., Kantartzis, K., Okada, H., & Kita, S. (2015). Sound symbolism facilitates word learning in 14-month-olds. *PLoS One*, 10(2), e0116494.
- Jones, M. N., Kintsch, W., & Mewhort, D. J. (2006). High-dimensional semantic space accounts of priming. *Journal of Memory and Language*, 55(4), 534-552.
- Joulin, A., Grave, E., Bojanowski, P., Douze, M., Jégou, H., & Mikolov, T. (2016). Fasttext. zip: Compressing text classification models. arXiv preprint arXiv:1612.03651.
- Keuleers, E. (2013). vwr: Useful functions for visual word recognition research. R package version 0.3. 0.
- Keuleers, E., & Brysbaert, M. (2010). Wuggy: A multilingual pseudoword generator. *Behavior Research Methods*, 42(3), 627-633.
- Köhler, W. (1929). *Gestalt psychology*. Liveright Publishing Corporation.
- Kounios, J., & Holcomb, P.J. (1994). Concreteness effects in semantic processing: ERP evidence supporting dual-coding theory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20(4), 804.
- Kruskal, J.B. (1968). Transformations of data. In D.L. Sills (Ed.), *International Encyclopedia of the Social Sciences*. Macmillan & Free Press.
- Küper, K., & Heil, M. (2009). Electrophysiology reveals semantic priming at a short SOA irrespective of depth of prime processing. *Neuroscience Letters*, 453(2), 107-111.
- 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.
- 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.
- 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. (2017). Word-Embeddings Italian Semantic Spaces: a semantic model for psycholinguistic research. *Psihologija*, 50(4), 503-520.
- Marelli, M., & Baroni, M. (2015). Affixation in semantic space: Modeling morpheme meanings with compositional distributional semantics. *Psychological Review*, 122(3), 485.

- Marelli, M., Traficante, D., & Burani, C. (2020). Reading morphologically complex words: Experimental evidence and learning models. In *Word Knowledge and Word Usage* (pp. 553-592). De Gruyter Mouton.
- Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. Retrieved from <https://arxiv.org/abs/1301.3781>.
- 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.
- Monaghan, P., Mattock, K., & Walker, P. (2012). The role of sound symbolism in language learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 38(5), 1152.
- Montefinese, M., Ambrosini, E., Fairfield, B., & Mammarella, N. (2014). The adaptation of the affective norms for English words (ANEW) for Italian. *Behavior Research Methods*, 46(3), 887-903.
- Nagy, W. E., & Anderson, R. C. (1984). How many words are there in printed school English?. *Reading Research Quarterly*, 304-330.
- Nation, I. (2006). How large a vocabulary is needed for reading and listening?. *Canadian Modern Language Review*, 63(1), 59-82.
- 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.W., Gray, J.R., Simpson, S., MacAskill, M.R., Höchenberger, R., Sogo, H., Kastman, E., Lindeløv, J. (2019). PsychoPy2: experiments in behavior made easy. *Behavior Research Methods*, 51(1), 195-203.
- Perniss, P., & Vigliocco, G. (2014). The bridge of iconicity: from a world of experience to the experience of language. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1651), 20130300.
- 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.
- Petilli, M. A., Günther, F., Vergallito, A., Ciapparelli, M., & Marelli, M. (2021). Data-driven computational models reveal perceptual simulation in word processing. *Journal of Memory and Language*, 117, 104194.
- 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.
- Ramscar, M. (2019). Source codes in human communication. arXiv preprint arXiv:1904.03991.

- 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*. RStudio, Inc., Boston, MA URL <http://www.rstudio.com/>.
- Rueckl, J. G., & Olds, E. M. (1993). When pseudowords acquire meaning: Effect of semantic associations on pseudoword repetition priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 19(3), 515-527.
- Saffran, J. R., & Wilson, D. P. (2003). From syllables to syntax: multilevel statistical learning by 12-month-old infants. *Infancy*, 4(2), 273-284.
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, 274(5294), 1926-1928.
- Sapir, E. (1929). A study in phonetic symbolism. *Journal of Experimental Psychology*, 12(3), 225.
- Seidenberg, M. S., & Gonnerman, L. M. (2000). Explaining derivational morphology as the convergence of codes. *Trends in Cognitive Sciences*, 4(9), 353-361.
- Stark, C.E., & McClelland, J.L. (2000). Repetition priming of words, pseudowords, and nonwords. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26(4), 945.
- Thompson, S. P., & Newport, E. L. (2007). Statistical learning of syntax: The role of transitional probability. *Language Learning and Development*, 3(1), 1-42.
- 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(3), 275-296.
- Van der Loo, M. P. (2014). The stringdist package for approximate string matching. *The R Journal*, 6(1), 111-122.
- 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.
- Xavier Alario, F., Segui, J., & Ferrand, L. (2000). Semantic and associative priming in picture naming. *The Quarterly Journal of Experimental Psychology*, 53(3), 741-764.
- Yarkoni, T., Balota, D., & Yap, M. (2008). Moving beyond Coltheart's N: A new measure of orthographic similarity. *Psychonomic Bulletin & Review*, 15(5), 971-979.

Table 1

Example of the sub-word vectors retrieved in order to represent the vector of the word <eating>.

Word	Length(n)	Character n-grams
asking	3	<as, ask, ski, kin, ing, ng>
asking	4	<ask, aski, skin, king, ing>
asking	5	<aski, askin, sking, king>
asking	6	<askin, asking, sking>

Figure 1

Correlation matrix between the continuous predictors included in the re-analysis of the SPP data for the short (A) and long SOA (B). Asterisks indicate statistically significant correlations.

Figure 2

Results of SPP data on short (A) and long SOA (B) on RTs, and short (C) and long (D) SOA on accuracy. SRel had an opposite effect on RTs and accuracy depending on whether the type of pair was word-word (red) or word-pseudoword (blue), enhancing word and impairing pseudoword recognitions in the short SOA, while the results in the long SOA were less consistent.

Figure 3
Correlation matrix between the continuous predictors included in Experiment 2. Asterisks indicate statistically significant correlations.

† Note: the label "Target-related frequency" refers to two different indexes based on target type.
For target words it is their frequency; for target pseudowords it is the frequency of the first orthographic neighbor.

Figure 4

RTs (A) and accuracy (B) results of Experiment 2. For the word-word pairs (red), the higher the SRel, the faster the RTs and the higher the accuracy; on the contrary for the word-pseudoword pairs (blue), the higher the SRel, the slower the RTs and the lower the accuracy.