LANGUAGE, COGNITION AND NEUROSCIENCE
https://doi.org/10.1080/23273798.2025.2610245

Routledge
Taylor & Francis Group

OPEN ACCESS Check for updates

Does morphology support novel word learning through statistical learning?

Olga Solaja[a], Jon W. Carr[a] and Davide Crepaldi[a,b]

[a]Cognitive Neuroscience Area, International School for Advanced Studies (SISSA), Trieste, Italy; [b]Department of Nervous System and Behavior Sciences, University of Pavia, Pavia, Italy

ABSTRACT

Most novel words that speakers learn are morphologically complex (e.g. *columnist*, *whistleblower*). Rather than hindering learning, such complexity might provide facilitation. Affixes (e.g. *pre-*, *-ness*) could do so in at least two ways: either thanks to their meaning or because they are frequent clusters of letters. This information is statistical in nature and might thus be captured via statistical learning. To investigate this, Italian native speakers learned novel words with (i) existing suffixes (*rugob-enza*, akin to *troft-er* in English), (ii) non-meaningful endings matched in frequency (*rugob-ondo*, *troft-an*), and (iii) non-meaningful, low-frequency endings (*rugob-allo*, *troft-ov*). Participants also completed a visual statistical learning task. Results showed that items with suffixes and low-frequency endings were learned best, while also exhibiting the strongest correlation with statistical learning ability. These results highlight the role of statistical information in word learning, and particularly its connection with the effect of derivational morphology in the process.

ARTICLE HISTORY
Received 10 January 2025
Accepted 12 December 2025

KEYWORDS
Word learning; statistical learning; morphology; reading; visual word recognition

Introduction

Compositionality is a key feature of human language, which allows us to convey a possibly infinite set of meanings based on a limited number of fundamental units. Derivational morphology is a core mechanism in this process. For example, the suffix – *able* indicates the capacity to be a certain way (e.g. *usable*, *likeable*). This process of constructing new lexical items is highly efficient, as understanding the meaning of the verb (*use*) and the suffix (*-able*) provides a complete understanding of the derived word (*usable*), even if the specific combination has not been encountered before.

Morphology, and in particular derivational morphology (i.e. prefixation, suffixation), has been extensively studied within the field of visual word recognition, where research has established morpheme recognition as a crucial step in the reading process (for a review, see Amenta & Crepaldi, 2012). Most of this evidence comes from lexical decision studies. For example, it has been consistently demonstrated that the time taken to recognise a complex word is directly related to the frequency of occurrence of its stem (e.g. New et al., 2004). In skilled readers, morpheme identification is automatised to such an extent that people decompose complex words into their constituent morphemes even outside of awareness, as shown in masked priming studies (e.g. Longtin et al., 2003; Rastle et al., 2004). Due to their semantic role, morphemes occur frequently across different words, which makes them recurring clusters of letters (e.g. – *able* is a much more frequent letter cluster than – *oble* or – *afle*). Their twofold nature as meaning-bearing and recurring orthographic units presents an intriguing contrast with other letter clusters that are equally frequent in language, but do not carry meaning (e.g. *con-*, as in *convert*, *contain*, *confetti*, etc.).

Determining whether morphological effects arise from word frequency or meaning is difficult, since these factors are closely intertwined in natural language. Nonetheless, research has started to examine the specific role of orthographic factors in morphological effects. Longtin et al. (2003) and Rastle et al. (2004) found that words that have only an apparent morphological structure (e.g. *corn-er*, which resembles *farm-er* but lacks real morphological meaning) can prime the recognition of embedded “stems” (e.g. *corn*), which suggests that early morphological parsing occurs regardless of semantic transparency. This behaviour mimics the effects observed with genuinely

CONTACT Olga Solaja olga.solaja@gmail.com Cognitive Neuroscience Area, International School for Advanced Studies, Via Bonomea 265, 34136, Trieste, Italy

Supplemental data for this article can be accessed online at https://doi.org/10.1080/23273798.2025.2610245.

morphologically complex words (*farmer* primes *farm*). Such evidence strongly suggests that segmentation into morphemes can occur independently of semantic factors (Amenta et al., 2017; Hasenäcker et al., 2016; Marelli et al., 2015). Of course, this effect reflects word processing, as attested by the fact that the *corner-corn* literature is primarily (if not exclusively) based on lexical decision/masked priming studies. Though these effects might obviously depend on how readers learn about morphemes, the core point here is that in early processing at least, they tend to identify morphemes even in words where they don't carry meaning (but see, e.g. Feldman et al., 2009, and Feldman et al., 2012, for results that are taken to support a broader role for semantics).

Notably, these effects depend on the morpheme position within the word. Crepaldi et al. (2010b) demonstrated that suffixes are processed as morphemes only when they appear in their typical position. For example, the nonword *gasful* produces interference in a lexical decision task, while *fulgas* does not. This body of evidence has fostered several competing models of complex word processing (Beyersmann & Grainger, 2023; Crepaldi et al., 2010a, 2010b; Grainger & Ziegler, 2011; Taft & Nguyen-Hoan, 2010), which collectively converge on the view that readers automatically segment printed words into morphemes early in processing – independently of meaning – and that this segmentation is sensitive to morpheme position. These findings clarify how morphological structure supports the rapid identification of printed, familiar words.

Importantly, sensitivity to morphemes extends to nonwords as well. For example, in an experiment carried out in Italian by Burani et al. (2002), nonwords made up of stems and suffixes (e.g. *woman-ist*) were categorised more frequently as possible words and were named more quickly and accurately than matching nonwords without suffixes (e.g. *woman-ost*). Similar findings emerged in several other languages (e.g. Spanish: Duñabeitia et al., 2008; English: Taft & Forster, 1975). However, from this body of literature alone, it is unclear whether this information is utilised during word learning; none of these experiments specifically explored the acquisition of unfamiliar stimuli as novel, meaningful lexical items.

A step in this direction has been made by studies explicitly examining the learning of morphologically complex novel words (e.g. Dawson et al., 2021; Ginestet et al., 2021; Havas et al., 2015; Merkx et al., 2011; Solaja & Crepaldi, 2024; Tamminen et al., 2015). For example, Ginestet et al. (2021) studied the acquisition of nonwords composed of novel, unfamiliar stems and existing prefixes and suffixes (e.g. *re-lerb-er*) in an eye tracking paradigm. After reading these novel words in short narrative paragraphs, participants showed more accurate spelling for the complex words than for orthographic controls. In this particular study, the role of derivational morphology was less clear in the eye tracking data; for example, there was a significant effect of word structure on relatively later metrics of lexical processing (gaze duration), but not on earlier eye tracking indicators (single fixations).

In another eye tracking study with Italian readers, Solaja and Crepaldi (2024) exposed their participants to complex novel words embedded in otherwise meaningful and familiar sentences (e.g. "John typically eats *pibs* for lunch".). Their eye tracking data clearly showed that reading novel suffixed items (e.g. *rugob-enza*; Eng. *troft-er*) is easier compared to noncomplex items (e.g. *rugob-ondo*; Eng. *troft-an*). Moreover, they found that the suffixed items were more accurately recognised in a verification memory task.

Tamminen et al. (2015) explored how participants acquire novel words composed of familiar stems and new suffixes (e.g. *crab-afe*). Definitions for the novel words were created such that the meaning of each novel affix modified the stem in a consistent way. Two key findings emerged: firstly, participants successfully extracted the meaning of the suffixes and generalised this knowledge to untrained novel words after a period of memory consolidation (see also Beyersmann et al., 2021, 2023). Secondly, the results showed that some statistical properties of the constituents, such as morphological family size (i.e. how many stems an affix has been attached to) affected learning: items with a larger morphological family size yielded better performance than those with smaller families. Importantly, the advantage of a large morphological family has been replicated in a study by Behzadnia, Ziegler, et al. (2024), which compared the novel word learning in big vs. small morphological families. Along similar lines, Beyersmann et al. (2024) demonstrated that the frequency of individual constituents in novel morphologically complex words also influences learning. In this study, participants were exposed to complex items where both constituents were unfamiliar (e.g. *ansu-da*). The authors independently manipulated the frequency of the first and the second element in the novel words, and found a complex pattern of results that included both facilitatory effects (better processing for higher-frequency constituents) and inhibitory effects (better processing for lower-frequency constituents). While the former are easy to explain, Beyersmann et al. (2024) drew on the Rescorla-Wagner approach (Rescorla, 1988; Shanks, 1995; Wagner and Rescorla, 1972) and its *unlearning* principle (e.g. Ramscar et al., 2013) to account for the latter. This approach is fundamentally based on the association between cues and outcomes. When a cue is present and an outcome is also present, the two get associated. But when a cue is present and an outcome is not, the two events get *dissociated*; the association learned in the past is weakened (hence the *unlearning* principle). In this framework, a higher frequency of occurrence is not necessarily a blessing; as much as there are more chances to enter into a strong association (for example, when a given bit of orthography is systematically associated with a given bit of meaning, or two novel morphemes often appear together), there are also more chances to lose those associations when the input or the environment is more diverse and irregular. In essence, there is a trade-off between the advantage of many exposures and the diversity that, everything else being equal, comes with it. Take two suffixes like – *er* (as in *singer*) and – *ade* (as in *blockade*). The former is quite frequent as a word ending in general, and therefore you'd find several genuine suffixed words where one learns the association between – er and the agentive meaning (e.g. *dealer*, *seller*, *worker*, *teacher*), but also several words where the expectation of finding an agentive meaning given the presence of a final – *er* is broken (e.g. *darker*, *larger*, but also *corner*, *beer*, *deer*). From the perspective of Rescorla-Wagner associative learning, these latter instances would substantially hamper the link between – *er* and an agentive meaning, therefore making any effect related to the presence of the suffix – *er* weaker. A relatively rare word ending would exercise a much smaller such effect, merely by the fact that there are much less words ending with that chunk, and therefore less occasions of an "unlearning" event.

Overall, the evidence illustrated above shows that morphological information does influence the learning of novel words, even without explicit instruction about morphology. There is also clear evidence that the statistical properties of the novel words' constituents play an important role. However, there doesn't seem to be a clear picture yet as to the specific mechanisms that might underlie these effects. Also, it is not clear what is the *source* of this sensitivity to statistical information in novel word learning – whether this is a reflection of a more general ability to capture statistical information in the perceptual input (e.g. Fiser and Aslin, 2001; Frost et al., 2019; Saffran et al., 1996) or a domain-specific mechanism shaped by the particular demands of language (e.g. Bogaerts et al., 2020; Siegelman et al., 2020).

Recent research from the field of statistical learning offers a complementary perspective on this issue. As mentioned above, derivational morphology establishes regularities not only in the mapping between form and meaning, but also in the way letters co-occur within words. Lelonkiewicz et al. (2020) trained participants using pseudoletter strings (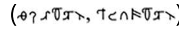), each containing an affix-like chunk () that occurred across several items. These affix-like chunks were positioned consistently within the string. During the testing phase, participants were presented with newly composed items, some of which included the trained affixes. Results showed that participants were more inclined to categorise newly composed testing items as belonging to the training set if they contained a trained affix. Because the experiment did not involve any orthographic, phonological, or semantic information, these data suggest that the extraction of morpheme-like chunks was based solely on visual statistical information, suggesting that there could be a role for statistical learning in structuring linguistic input. In a further study by the same group (Lelonkiewicz et al., 2023), linguistic information was progressively factored in (e.g. by using real letters instead of pseudocharacters, or associating strings with objects), and the learning effect increased in size.

Interestingly, the connection between reading, visual word identification, and our ability to discern the numerous regularities present in (written) languages has been explored beyond the domain of morphology. After all, letters might be statistically associated more generally, independently of whether they are part of morphemes. For instance, in English, upon encountering the letter C, we might naturally expect the letter A to follow, while the letter S is less likely. Chetail (2017) also used an alphabet unfamiliar to participants and embedded regularities (e.g. bigrams and their specific position within a letter string) in a set of novel words that participants were asked to learn. Participants were found to judge novel combinations of letters containing frequently occurring bigrams as more word-like compared to random letter combinations. Similar to Lelonkiewicz et al. (2020), this research removed any semantic, orthographic, or phonological information, yet the observed effects align with those typically associated with orthographic processing.

There have also been efforts to study more directly the contribution of statistical learning to reading skills. One promising avenue in this respect is the study of individual differences (e.g. Siegelman, 2020). This approach offers clear predictions at the individual level: if statistical learning is a relevant mechanism in language acquisition, individuals who excel in statistical learning should also excel in language learning. For instance, Siegelman et al. (2020) predicted that early readers who are more adept at learning statistical regularities would exhibit stronger reading abilities compared to those with weaker statistical learning (SL) abilities. The results indeed showed that readers who were more sensitive to orthography – phonology correspondences in a word naming task were better readers compared to those who relied more on orthography – meaning relationships. It is important to note that the strength of the correlation between statistical learning and reading skills varies across different studies and could be influenced by the specific task or measure used to assess reading abilities (e.g. Schmalz et al., 2019; Siegelman et al., 2017). Recent work (Siegelman et al., 2017; Siegelman & Frost, 2015) has pointed to the fact that the tests used in a large part of the published literature yield reliable effects at the group level, but not at the individual level. This has been addressed by developing a classic visual statistical task with improved psychometric properties (Siegelman et al., 2017). In addition to providing a rather wide distribution of participants' scores – like one would want from a task focused on individual variability rather than group differences – this task showed improved internal reliability and consistency.

This body of evidence, we believe, yields good support to the idea that readers acquire substantial knowledge of the statistical regularities encoded in the written language, and this supports their processing of letters and words. However, these results do not shed enough light on the mechanisms that lie behind the connection between statistical learning and reading, as attested most clearly by the opposite directions of some effects (e.g. letter chunk frequency in Beyersmann et al., 2024) and the volatility of correlational data.

Overall, morphology clearly plays a very important role in visual word identification and has been shown to provide a valuable anchor when we process unfamiliar, novel words. It is not clear, however, what precise role morphemes play in word learning, and what is the distinct contribution of letter chunk frequency on one hand, compared to the associations of these chunks with a particular meaning on the other. Because these phenomena are all related to linguistic regularities, statistical learning might be a powerful contributing factor. With this picture in mind, we investigated the learning of novel, artificial lexical items formed by a nonexisting stem (e.g. *troft-*) and: (i) an existing suffix (*troft-er*); (ii) a letter cluster of similar frequency, but with no association with meaning (*troft-an*); or (iii) a letter cluster with a lower frequency of occurrence (*troft-ov*). With such a design, we can tease apart the contributions of meaning and frequency: any difference between items (i) and (ii) can only be due to meaning, while any difference between items (ii) and (iii) can only be due to frequency. To obtain evidence on the potential connection between the detection of statistical regularities and the role of morphology in word learning, we complemented this design with a statistical learning task aimed at assessing the general statistical learning skills of our participants (Siegelman et al., 2017), which we correlated with their performance in the three conditions in the main learning task.

The study was preregistered at https://aspredicted.org/WBL_QSF, and included the following predictions. First, we predict that in the novel word learning task, high-frequency endings words would be learned better compared to the low-frequency endings words. The underlying hypothesis is that learners will recognise the highly frequent chunk as more familiar, which would lead them to better remember the whole item. Second, we predicted that suffixed words will exhibit better performance than high-frequency endings words due to the strong form-to-meaning relationship inherent to suffixes. We hypothesised that beside the familiarity with the visual aspect of the letter chunk (which is also present in the high-frequency endings case), the semantic content of the suffixes will provide a strong anchor for memory. Finally, we predicted that a correlation might emerge between statistical learning skills and novel word learning in the conditions where statistical regularities are present – that is, with suffixed items and, perhaps to a lesser extent, high-frequency endings items.[1] The rationale is that suffixes might benefit from two types of statistics: letter combinations frequency and the strength of their association with the meaning. On the other hand, high-frequency endings feature only one statistic. Therefore, participants skilled in statistical learning might have more opportunity to show an advantage in cases where there are more statistical regularities.

The data and code are available at https://osf.io/9kt74/?view_only=9861128b316e4f65b642de18ea575221.

Methods

The experiment was programmed using jsPsych (de Leeuw et al., 2023; version 7.1.2) and was hosted on the Cognition.run platform (www.cognition.run).

Participants

Two hundred and six participants were recruited through Prolific (www.prolific.co). The sample size was computed based on data simulations (see below). All participants were native speakers of Italian with normal or corrected-to-normal vision and no reading disabilities. The Prolific filtering system was used to limit the study to this population. They were paid for their participation a base amount of 5 GBP and received an additional bonus of up to 1.14 GBP according to performance (0.01 GBP for each correct answer). The experiment was approved by the SISSA Ethics committee (protocol number 326) and participants gave their informed consent before the start of the experiment.

Procedure

The experiment consisted of two main parts. The first part tested novel word learning and the second assessed general statistical learning skills. Participants received instructions at the beginning of each part and were informed of the bonus they may earn.

Novel word learning task

Stimuli

Participants were exposed to novel words that consisted of a stem and an ending. All stems were nonexisting words in Italian without any lexical neighbour (e.g. *rugob-*; an English example would be *troft-*). Specifically, they were either five or six characters long; their mean OLD20 (Yarkoni et al., 2008) and log bigram frequency were 2.09 (SD = 0.19) and 5.96 (SD = 0.30) respectively. These measures were calculated based on the SUBTLEX-IT word frequency database (Crepaldi et al., 2016). The stems were paired with three types of endings: (i) suffixes (e.g. – *enza*, token frequency mean = 4.77, SD = 0.49), (ii) nonmorphological endings matched to suffixes in token frequency in word-final position (e.g. – *ondo*, mean = 4.96, SD = 0.22), or (iii) nonmorphological endings with low frequency in word-final position (e.g. – *espa*, mean = 2.49, SD = 0.29). The endings were either three or four characters long and were all equally plausible phonotactically.

We used 72 stems and 24 endings, eight of each type. To create the novel words, we pseudorandomly paired each ending with three stems.[2] Thus, the final set of words to be learned consisted of 72 items, eight to ten letters in length (mean = 9.44) and with no lexical neighbours in Italian. We also made sure that none of the items contained an embedded existing Italian word. All participants were exposed to all the stimuli (i.e. there were no rotations).

The complete list of stimuli is reported in Appendix 1.

Procedure

During the learning phase of the task, a novel word appeared at the centre of the screen and participants were instructed to read it and remember it. In order to facilitate learning, they had to type it back into a textbox that appeared after a button press. Each word was repeated in a pseudorandom order three times during the learning, so each participant was exposed to 216 learning trials in total. The entire task was self-paced. Participants were informed that they would be tested on the novel words.

The learning phase was followed by a two-alternative forced choice task in which a trained item and its foil appeared together at the centre of the screen. The foils were created by substituting one of the letters in the stem (e.g. *rugobenza* vs. *rufobenza*), with the same constraints described above. We created two lists to counterbalance the position (i.e. whether the foil was on the left or the right side of the screen). There were a total of 72 trials, presented in random order. Participants were instructed to indicate which of the two items they recognised from the learning phase by pressing a marked key on the left or right.

Visual statistical learning task

This task followed closely the visual statistical learning task described by Siegelman et al. (2017). The only difference was the language of the instructions and the fact that our experiment was run online rather than in a laboratory setting. Notably, participants' statistical learning ability is assessed based on performance across a range of transitional probability difficulties and two task types, which together provide a more robust and reliable measure of statistical learning than is traditionally adopted.

Stimuli

The materials consisted of 16 shapes (Fig. 1), which were used to create eight triplets, further subdivided into four triplets with transitional probabilities (TP) between the shapes equal to 1, and four triplets with $TP = 0.33$. The

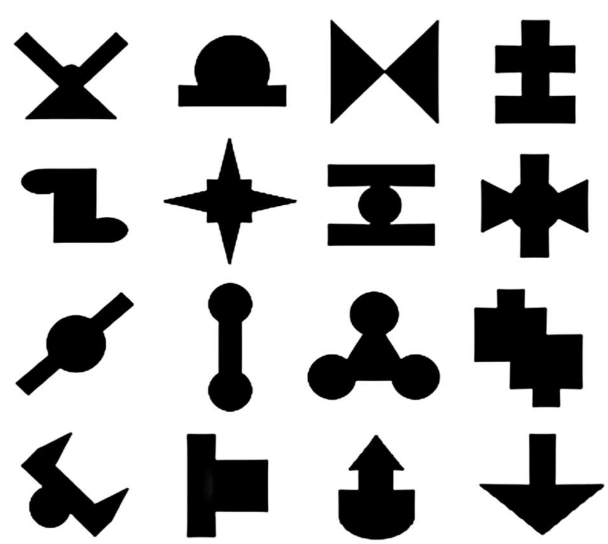

Figure 1. Shapes from the exposure phase.

combinations of shapes that constitute the triplets were randomised for each participant.

Procedure

In the exposure phase, the shapes were presented individually on the screen for 800 ms, with a 200 ms pause between stimulus presentations. Their sequence was determined by the triplet structure described above. Each of the eight triplets appeared 24 times during this phase, with the constraint that the same triplet could not appear two times in a row. The order was randomised in each run. Overall, the only cue to the structure of the stream were the TPs between the shapes. Participants were instructed to look at the sequence of shapes that appeared on the screen and were told that they would be tested on them. The testing phase included two different tasks. In one, participants were required to recognise a pattern from the training among two, three, or four distractors (34 trials; see Figure 2a). They were instructed to select the sequence that looked most familiar to them. In the second task, participants were presented with either one or two shapes from a triplet and had to select the missing shape from three alternatives (8 trials; see Figure 2b). They were instructed to select the shape that best completes the sequence. The statistical learning ability of each participant was computed as the total number of correct trials in these two tasks (total N = 42).

Power simulation

A power simulation was conducted to determine the minimum justified sample size. The methodology described in the study by DeBruine and Barr (2021) was employed, which involved running data simulations. In these simulations, we generated 1000 datasets where the effect sizes in the learning task were set to a 6% accuracy advantage for the high-frequency condition compared to the low-frequency condition. Additionally, there was an extra 6% advantage for the suffixed condition compared to the high-frequency condition, resulting in a total 12% accuracy boost for the suffixed condition compared to the low-frequency condition. The correlation between visual statistical learning (VSL) and the suffixed condition, as well as VSL and the high-frequency condition, was set to 0.4. In the absence of substantial literature to leverage, all the values above were informed by considerations on the minimal effect that we deemed of theoretical interest. The noise parameters are outlined in Table 1. Using this model, we achieved 100% power for the learning task and 79.9% power for the correlations with a sample size of 160 participants.

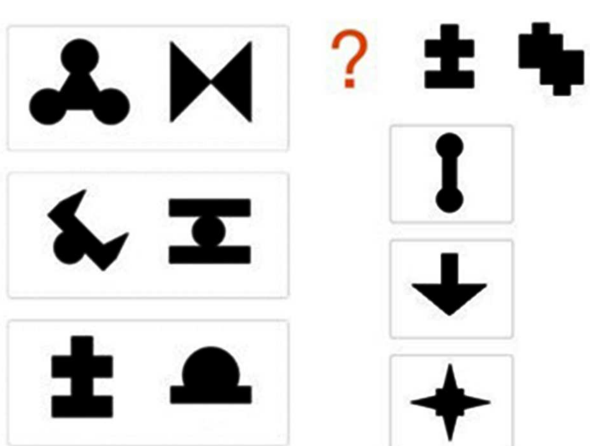

Figure 2. (a) Example of a trial in Pattern recognition task with two distractors. Instructions: "Please choose pattern you are most familiar as a whole" (b) Example of a trial in Pattern completion task with two distractors.

Results

The data were analyzed using the R programming language (R Core Team, 2025) using *tidyverse* (Wickham et al., 2019), *data.table* (Dowle & Srinivasan, 2023), *lmerTest* (Kuznetsova et al., 2017), *effects* (Fox, 2003), and *ggplot2* with associated packages for visualisation (*wesanderson*, *ggbeeswarm*, *ggsignif*, *cowplot*, *corrplot*, *ragg*) (Ahlmann-Eltze & Patil, 2019; Clarke & Sherrill-Mix, 2017; Pedersen et al., 2023; Ram et al., 2019; Wei et al., 2021; Wickham, 2016; Wilke, 2020).

Participants' exclusion

As per the preregistration, we excluded participants based on their performance in the typing task during the training phase. We calculated the total Levenshtein edit distance between the target words and the words that participants typed back; participants were excluded from the analysis when this index was higher than 30. As a result, we excluded 46 participants, which left 160 participants for the final analyses (male: 86, mean age: 28.81, range: 19–60 years).

Novel word learning task – 2AFC

The overall mean accuracy in the novel word learning task was 0.78 (SD = 0.41). Figure 3 breaks down test

Table 1. Parameters defining noise in the simulation model.

Overall	Item	Subject intercept learning task	Subject frequency slope	Subject meaning slope	Subject intercept VSL task
0.06	0.03	0.06	0.03	0.03	0.8

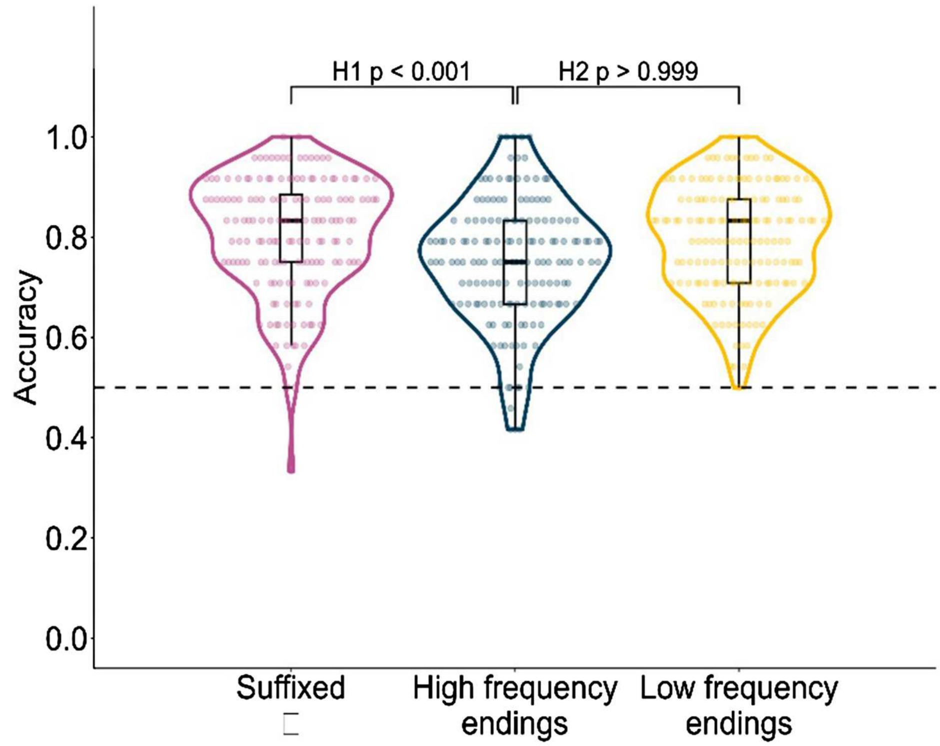

Figure 3. Accuracy in Novel word learning task across the ending types.

accuracy by condition, while Figure 4 shows mean difference between conditions. As predicted, and as per the preregistered analysis, suffixed words were learned better than high-frequency ending words (0.81 ± 0.12 vs. 0.75 ± 0.12; $t[159] = 6.21$, $p < .001$). This contrast unveils the effect of the ending having a meaning, while holding frequency constant. We also expected that high-frequency ending items would be learned better than their low-frequency ending counterparts. This prediction was not confirmed (0.75 ± 0.12 vs. 0.80 ± 0.11; $t[159] = -5.21$, $p > .999$). In fact, when we ran a two-tailed t-test on the same comparison, which was not preregistered, the low-frequency ending items turned out to be learned significantly better than the high-frequency ending items ($t = -5.21$, $df = 159$, $p < 0.001$), contrary to our hypothesis. We also ran a generalised mixed model as an additional, non-preregistered analysis which confirmed our results (see Supplemental material).

Visual statistical learning

The overall mean accuracy in the VSL test was 0.53 (SD = 0.50). The average number of correct responses by subject was 22.32 (SD = 5.54). The distribution of the scores can be seen in Figure 5.

Following our preregistration, we correlated the accuracy scores on the VSL task with the accuracy scores in each of the novel word learning conditions. As can be seen in Figure 6, the suffixed and low-frequency conditions showed a similar correlation with the VSL performance. With suffixed items, the correlation was 0.15, sitting just outside the conventional significance threshold ($t = 1.94$, $df = 158$, $p = 0.06$), while with low-frequency endings, the correlation was 0.18, below the significance threshold ($t = 2.32$, $df = 158$, $p = 0.02$). On the other hand, the correlation with accuracy in the high-frequency ending condition was smaller ($r = 0.05$) and not significant ($t = 0.64$, $df = 158$, $p = 0.53$). Again, we ran a non-preregistered generalised mixed model as a support to the correlation analysis, which found the same pattern (see Supplemental material).

Discussion

This work sits at the intersection between two lines of research. On the one hand, it is unclear what the specific role of derivational morphology is in the acquisition of novel words, particularly words that feature a known suffix and novel stem. On the other hand, recent studies (Chetail, 2017; Lelonkiewicz et al., 2020, 2023) have suggested a connection between certain morphological effects and statistical learning, specifically the frequency of letter chunks and letter co-occurrence. This leads to the prediction that individuals with good statistical learning abilities might be more

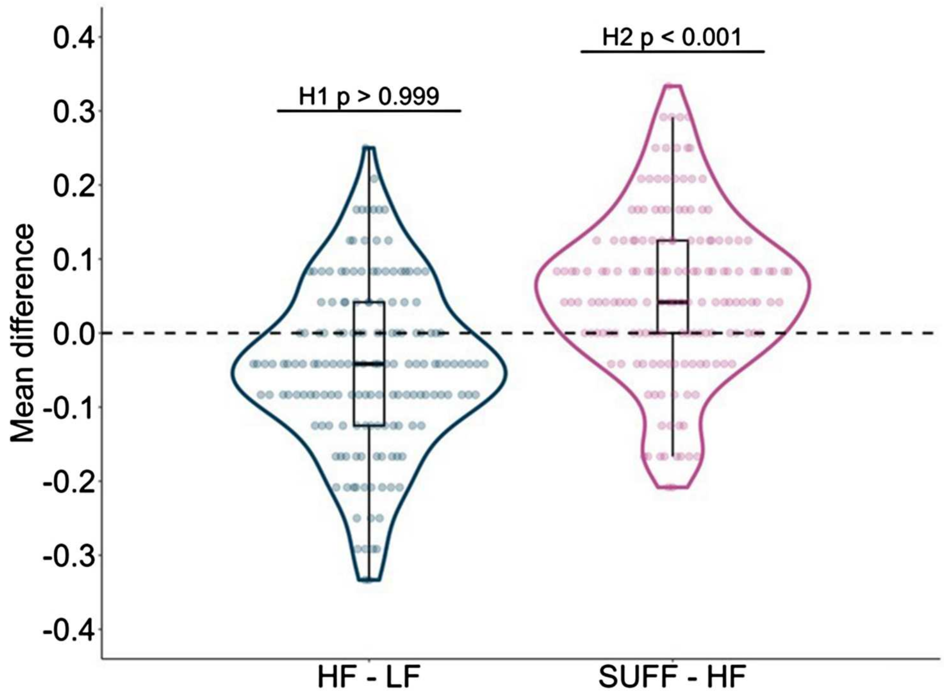

Figure 4. Mean difference between high-frequency vs. low-frequency endings, and suffixes vs. high frequency endings.

capable of learning novel words comprised of frequent letter clusters. In addition to that, they might also be more capable of using morphology in this process, when available, as an additional source of information that is ultimately based on statistical patterns.

To address these questions, we administered a novel word learning task where the items consisted of a nonword stem and a final chunk that could be a suffix (e.g. – *er*), an equally frequent, but non-meaningful ending in Italian (e.g. – *an*) or an infrequent ending (e.g. – *ow*). This design enabled us to specifically single out the effect of the frequency of the final chunk (comparing words with low-frequency vs. high-frequency endings), and the effect of meaning (comparing high-frequency endings vs. suffixed words). To investigate whether the possible frequency and meaning effects were connected to statistical learning, we tested our participants for their SL skills using the task developed by Siegelman et al. (2017) which was designed to capture individual variability.

Three main results emerged. First, we confirmed our predictions regarding the semantic effect: novel words with suffixes were learned more effectively compared to their counterparts with meaningless, high-frequency endings. Second, a frequency effect emerged, but it contradicted our initial expectations: items with low-frequency endings were better remembered than novel words with high-frequency endings. Third, the correlations between statistical learning skills and the word learning performance in the different conditions were

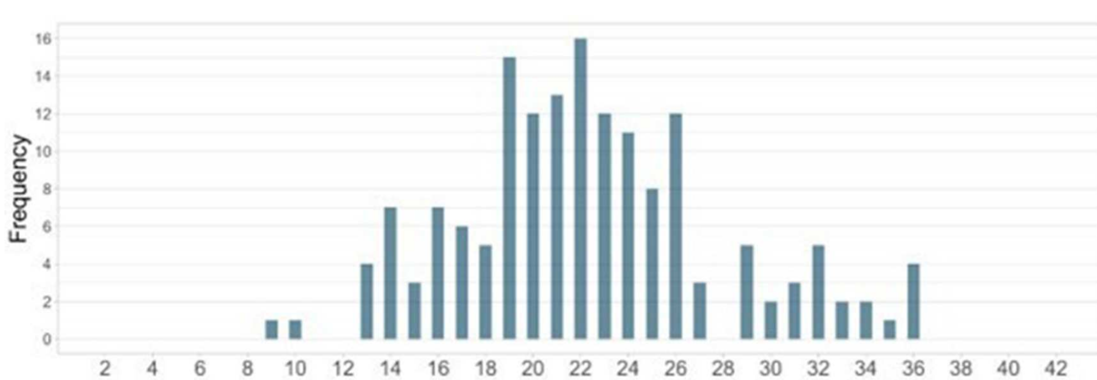

Figure 5. Visual statistical learning task. Distribution of scores.

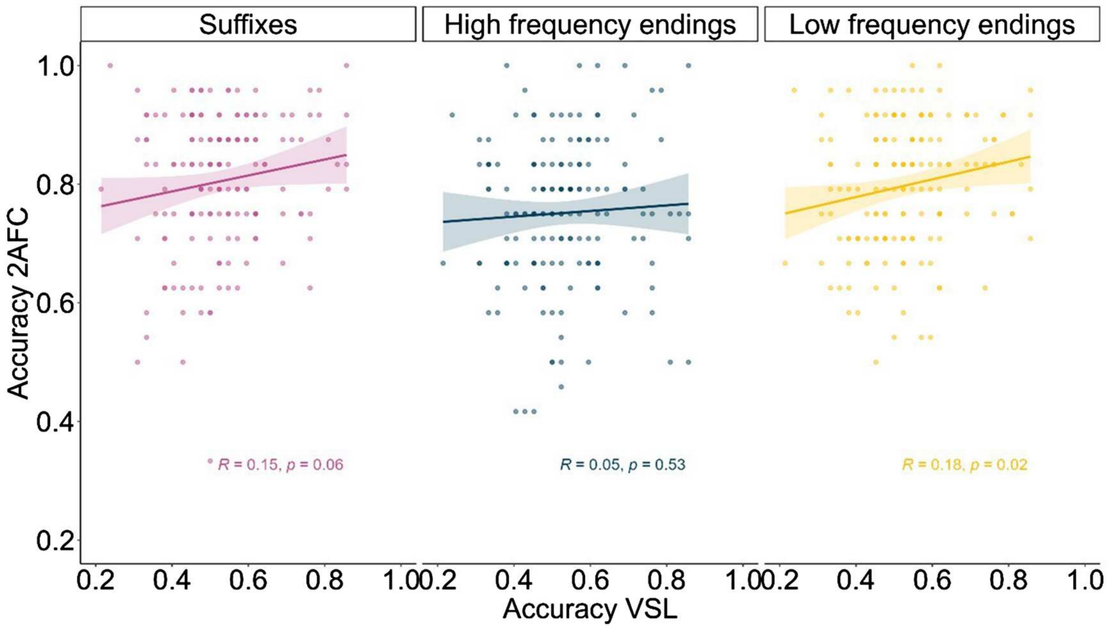

Figure 6. Correlation of Novel word learning and Visual statistical learning accuracy across three ending types.

not particularly strong, but lay around the significance threshold in the conditions where the best learning outcomes were observed (i.e. suffixed and low-frequency ending words). For the high-frequency ending condition, the correlation was not significant.

The suffix advantage confirms that derivational morphology plays a role in the learning of novel words (e.g. Dawson et al., 2021). Because this effect emerged in the comparison between suffixes and frequency-matched, non-meaningful items, this finding seems to be due to the fact that suffixes carry meaning, not to their frequency as letter chunks. These results extend the previous literature highlighting the role of morphology in the processing of novel words (e.g. Burani et al., 2002, 2008; Crepaldi et al., 2010b; Taft & Forster, 1975; Yablonski & Ben-Shachar, 2016; see Amenta & Crepaldi, 2012 for a review). Apparently, not only do we detect the morphological structure of unfamiliar letter strings, but we also actively use that information to store novel words in memory.

There are two possible explanations for these results. One is that novel words are learned more effectively due to the dynamics of the lexical system – that is, the underlying processes and patterns of activation that govern how stored linguistic representations, such as suffixes, are accessed and integrated during word learning and processing. According to several models of complex word identification (Crepaldi et al., 2010b; Grainger & Ziegler, 2011; Marelli & Baroni, 2015; Taft & Nguyen-Hoan, 2010), suffixes have distinct representations within the lexical system (while this is not the case for high-frequency, but meaningless letter chunks). When encountering a novel word, these representations are activated, facilitating the formation of a lexical representation for the novel word. Importantly, this account refers to the same lexical dynamics that are generally believed to underlie the processing of familiar words (Coltheart et al., 2001; McClelland & Rumelhart, 1981) – dynamics that are automatic, implicit, and operate independently of strategic control or conscious awareness.

Another possibility is that familiar suffixes contribute to the improvement of episodic memory for the novel words, resulting in better retention and recall of these words. This might again be related to the structure of the underlying lexical system, of course. However, in this account, this phenomenon *per se* would primarily relate to the workings of explicit memory, rather than to the implicit, non-conscious dynamics of the lexical system. In other words, while demonstrating that suffixes do have a role in word learning that is distinctive and separate from other letter chunks, the present data open a new question: are we constructing a new implicit representation of the novel words specifically linked to the representation of the suffixes within them? Or rather, are the suffixes merely making us remember better that we have encountered a novel word in the recent past? Of course, the two possibilities are not necessarily mutually exclusive; we might have a better episodic memory for suffixed novel words, but perhaps this is precisely because those novel words activate a suffix representation in the lexical system. Future work might address this issue.

These theoretical considerations relate to the vast literature on the consolidation of novel lexical entries – the process by which newly encountered words are integrated into a person's long-term memory, and stored as representations that interact with previously existing ones (e.g. Davis & Gaskell, 2009; Plaut et al., 1996; Ullman, 2001). The morphological effect observed here unveils the possibility that suffix representations are involved in the process of lexical consolidation, which in turn opens a new set of interesting questions. How does this happen? Does this create new ties in the lexical system that, for example, might be amenable to investigation via completely implicit phenomena like priming (see, e.g. Viviani & Crepaldi, 2022)? And under which conditions would a suffix-mediated consolidation emerge? Does it depend on the previous status of the morphological representations, so that it could bring more benefit to one individual than another (such as in "semantic" vs. "orthographic" profile in Andrews and Lo (2013) study)? Does it require several encounters with the novel word, or a sufficiently varied experience (e.g. Mak et al., 2021; Tamminen et al., 2015)?

Our findings clearly suggest that the learning advantage provided by suffixes during the acquisition of novel words is due to the involvement of more linguistic and meaning-oriented morphological representations. This conclusion is supported by the fact that non-meaningful letter chunks, despite being equally frequent, did not provide the same learning advantage. Different models of morphological processing have different ways of addressing the contrast between more visual, frequency-based stages of processing on the one hand, and more linguistic, meaning-rich morphological processing (Beyersmann & Grainger, 2023; Crepaldi et al., 2010b; Grainger & Beyersmann, 2017; Grainger & Ziegler, 2011; Taft & Nguyen-Hoan, 2010). However, all these models agree that both types of representations/processing exist. Our results align with previous research that suggests that establishing links between form and meaning in new lexical items is key to learning such items (Dawson et al., 2021).

The second important finding of this experiment is that novel words with low-frequency endings were learned significantly better than their high-frequency counterparts. Frequency effects are very well established in the processing of familiar words (e.g. Longtin et al., 2003; New et al., 2004; Rastle et al., 2004), but from our results, it appears that frequency of occurrence yields opposite effects in novel word learning vs. familiar word identification. This finding, although contrary to our predictions, replicates the results of Beyersmann et al. (2024), which also found low-frequency endings to facilitate learning.

In line with the proposal of Beyersmann et al. (2024), these somewhat counterintuitive results could be explained in terms of the negative learning principle (also referred to as *unlearning*; Ramscar et al., 2013; Rescorla, 1988; Shanks, 1995). In essence, the frequency with which a particular input occurs contributes to both positivx`e learning – when that input is consistently paired with a specific output or event – and negative learning – when the same input is paired with multiple, different outputs or events. In the context of our experiment, unlearning would occur each time an ending is paired with a new stem, thus breaking the pattern coming from previous encounters of the same suffix with a different stem. This disruption would be more pronounced for high-frequency endings, which in real language occur more often, and therefore have been attached to many different stems.

Another possibility is that novel words with low-frequency patterns are more striking, which potentially made them more memorable. This aligns with the literature on word and sentence processing that focuses on concepts such as *surprisal* (e.g. Amenta et al., 2023; Frank & Bod, 2011; Hale, 2001; Levy, 2008). Amenta et al. (2023) showed that there is a tradeoff between sentence-level and word-level cues, and that reader behaviour is effectively described based on the amount of information carried by a given word in the context of a given sentence. Complementary to this, error-driven learning postulates that larger mismatches between expected and observed states generate stronger memories (e.g. Henson & Gagnepain, 2010). From this perspective, low-frequency endings may have triggered better learning by violating the general word-ending pattern in the artificial lexicon used in this study.

An interesting aspect of the results is that the performance in the suffixed and low-frequency conditions in the word learning task was nearly identical (80% vs. 81%, respectively). This suggests that the mechanisms described above are both at work together, even if they seem to capture different, or even opposite types of statistics. This in turn suggests that the cognitive system flexibly engages with different types of regularities – morphological structure in the case of suffixes, and statistical irregularity in the case of low-frequency endings – depending on what cues are most informative in the learning environment. Rather than reflecting two competing mechanisms, the system may instead deploy different strategies based on the available input. Of course, it is hard to generalise based on these findings only; we have focused here on two very specific types of regularities in the written language, among many. We hope that future work will extend the present investigation and provide converging evidence on a more varied set of statistics. It would also be useful, of course, to see a replication of the present findings, which appear to be statistically robust (based on a rather severe, a priori power analysis, for example), but were unexpected, at least in part.

The individual variability results shed more light on the overall theoretical message of these findings. Statistical learning skills exhibited correlation with the learning of suffixed and low-frequency-ending items, while no correlation was found with novel words containing high-frequency endings (these correlations were not very strong overall, however; more on this below). It is not difficult to see why statistical learning might be specifically related to morphology. As a symbolic system, the human language is characterised by a fundamental arbitrariness in the way form conveys meaning; there is generally nothing in the way a word sounds, or looks, that is suggestive of its meaning (e.g. Hockett, 1963; but see research on sound symbolism for cases where form does seem suggestive of meaning, e.g. Sapir, 1929; Imai & Kita, 2014). Derivational morphology, however, establishes some regularity in this mapping and thus introduces predictability. Although the association between, for example, the suffix – *er* and its meaning remains arbitrary, the fact that, for example, *driver*, *settler*, *buyer*, *dealer,* all denote someone who performs a specific action increases the predictability of lexical mappings. Therefore, when we encounter a novel word such as *libber*, while its exact meaning might be unknown, we can confidently assume that a *libber* is someone who *libs*. Substantial evidence supports our sensitivity to these regularities (e.g. Marelli et al., 2015). Moreover, recent studies (Lelonkiewicz et al., 2020, 2023) have provided direct causal evidence that morphology-like statistics can induce morphology-like effects, even when the stimuli are entirely unfamiliar to the readers and lack any reference to semantics and/or phonology.

Interestingly, the correlation results highlight that statistical learning not only captures the correspondence between form and meaning (as in the case of suffixes), but also the statistics of orthography per se (as in the case of low-frequency endings). In fact, the only difference between high-frequency and low-frequency ending items was frequency itself – neither type of word ending is connected to any specific piece of meaning in the language. Therefore, the frequency of letter chunks seems to be another relevant piece of information that statistical learning captures and that affects the acquisition of novel lexical material. It is not entirely clear why the correlation manifests itself with low-frequency, but not high-frequency endings, given that both the extremes of the frequency distribution must be based on some form of statistical learning. We can only speculate here; perhaps the fundamental mechanisms behind the appreciation of very frequent vs. very surprising events are different? This would reflect nicely the long-standing debate between confirmation-based vs. error-based learning, which highlight the importance of either predictable or surprising events. On the other hand, it would not be clear why only one of the two mechanisms would be specifically linked with word learning, as the present data suggest. Clearly, more research is needed here.

These observations are in line with previous work investigating the role of letter statistics in visual word identification and processing (Chetail, 2017; Lelonkiewicz et al., 2020; Seidenberg, 1987; Treiman et al., 2018). Moreover, it directs our attention to a critical question in the statistical learning literature: since the information that is used in detecting letter statistics and inferring meaning is different, is the cognitive phenomenon that underlies the two correlations one and the same? Or rather, are some neural circuits/cognitive processes at work at the interface between form and meaning, while other, separate neural circuits/cognitive processes capture letter statistics? Our data certainly cannot settle this question, although the fact that the same individual SL skill correlates with both the performance on suffixed words and the performance on low-frequency endings might be taken to suggest that the same fundamental statistical learning process may support word learning in both conditions (for a broader discussion of this issue, see Bogaerts et al. (2022)).

Finally, it is important to stress that although the relevant correlations discussed above are of theoretical interest, they are still quite small. One reason for this could be the different training methods in the statistical learning and word learning tasks. While in the latter participants were exposed to static stimuli without any time constraints, in the visual statistical learning training, they encountered fast-paced, dynamic stimuli. Furthermore, during word learning, participants were actively engaged (i.e. they had to type back the words they were learning), whereas the SL task simply required passive viewing. These differences are certainly non-negligible; however, they make the correlations that we observe here even more compelling and provide further support to the idea that statistical learning contributes to word learning, especially for what concerns form and form-to-meaning regularities. At the same time, the fact that these correlations are not big in size also suggests that the role of statistical learning in these mechanisms might be fairly limited, and effective word learning is likely to depend on a host of cognitive processes (e.g. attention, working memory, vocabulary size; Gathercole & Baddeley, 1993; Nation, 2006, 2008;

Ouellette, 2006; Smith et al., 2010) where statistical learning is only one of many factors.

Overall, these findings highlight the important role of statistical information in the acquisition of novel lexical information, and clarify its connection with the role of derivational morphology in word learning. More specifically, it seems that the consistent correspondence between orthographic (and potentially, phonological) information on the one hand, and meaning on the other, facilitates the acquisition of novel words. Frequency of occurrence, per se, generates a surprising effect whereby less frequent material seems to promote better learning. This last finding in particular, but also the pattern of results observed here more generally, offer an interesting perspective on the specific mechanisms through which statistical information affects lexical dynamics. More broadly, this work moves the spotlight from the more basic question of *whether* statistical learning and reading are connected (the answer appears to be yes), to the more intricate issue of *how* the two cognitive skills interact.

Notes

1. A reviewer noted that our formulation of this prediction in the pre-registration document was ambiguous, and could easily be interpreted as referred to the difference in performance between the suffixed and the high-frequency-ending conditions on the one hand, and the low-frequency condition on the other – rather than with the suffixed and high-frequency-ending conditions per se. We apologise for this mistake, and we thank the reviewer for encouraging a clarification of this prediction in the manuscript.
2. Unfortunately, the usual procedure with completely randomised pairings implemented in a Latin square design led to too many items that contained existing words and were therefore not acceptable as stimuli.

Acknowledgements

This project was funded by the European Union - Next Generation EU, Mission 4 Component 1, CUP G53D23003240006 (prog. Code 2022Z7425F, "Learning to use words: the statistical nuances of category bootstrapping", PI: Davide Crepaldi).

The authors thank Yamil Vidal for the fruitful discussions that inspired this study.

Disclosure Statement

No potential conflict of interest was reported by the author(s).

Funding

Financed by the European Union - Next Generation EU, Mission 4 Component 1, CUP G53D23003240006 (prog. code 2022Z7425F, "Learning to use words: the statistical nuances of category bootstrapping", PI: Davide Crepaldi.

Open Practices Statement

The analysis and exclusion criteria were preregistered at https://aspredicted.org/WBL_QSF. The stimuli, code and the data for all experiments are available at OSF repository https://osf.io/9kt74/?view_only=9861128b316e4f65b642de18ea575221.

References

Ahlmann-Eltze, C., & Patil, I. (2019). *ggsignif: Significance brackets for 'ggplot2'*. R package version 0.6.0. v.

Amenta, S., & Crepaldi, D. (2012). Morphological processing as we know it: An analytical review of morphological effects in visual word identification. *Frontiers in Psychology*, *3*(JUL), 1–12. https://doi.org/10.3389/fpsyg.2012.00232

Amenta, S., Hasenäcker, J., Crepaldi, D., & Marelli, M. (2023). Prediction at the intersection of sentence context and word form: Evidence from eye-movements and self-paced reading. *Psychonomic Bulletin & Review*, *30*(3), 1081–1092. https://doi.org/10.3758/s13423-022-02223-9

Amenta, S., Marelli, M., & Sulpizio, S. (2017). From sound to meaning: Phonology-to-Semantics mapping in visual word recognition. *Psychonomic Bulletin & Review*, *24*, 887–893. https://doi.org/10.3758/s13423-016-1152-0

Andrews, S., & Lo, S. (2013). Is morphological priming stronger for transparent than opaque words? It depends on individual differences in spelling and vocabulary. *Journal of Memory and Language*, *68*(3), 279–296. https://doi.org/10.1016/j.jml.2012.12.001

Behzadnia, A., Ziegler, J. C., Colenbrander, D., Bürki, A., & Beyersmann, E. (2024). The role of morphemic knowledge during novel word learning. *Quarterly Journal of Experimental Psychology*, *77*(8), 1620–1634. https://doi.org/10.1177/17470218231216369

Beyersmann, E., & Grainger, J. (2023). The role of embedded words and morphemes in reading. In D. Crepaldi (Ed.), *Linguistic morphology in the mind and brain* (1st ed., pp. 26–49). Routledge. https://doi.org/10.4324/9781003159759-3

Beyersmann, E., Grainger, J., Dufau, S., Fournet, C., & Ziegler, J. C. (2024). The effect of constituent frequency and distractor type on learning novel complex words. *Language, Cognition and Neuroscience*, *39*(2), 251–264. https://doi.org/10.1080/23273798.2023.2263590

Beyersmann, E., Wegener, S., Nation, K., Prokupzcuk, A., Wang, H. C., & Castles, A. (2021). Learning morphologically complex spoken words: Orthographic expectations of embedded stems are formed prior to print exposure. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *47*(1), 87–98. https://doi.org/10.1037/xlm0000808

Beyersmann, E., Wegener, S., Spencer, J., & Castles, A. (2023). Acquisition of orthographic forms via spoken complex word training. *Psychonomic Bulletin & Review*, *30*(2), 739–750. https://doi.org/10.3758/s13423-022-02185-y

Bogaerts, L., Frost, R., & Christiansen, M. H. (2020). Integrating statistical learning into cognitive science. *Journal of*

Memory and Language, *115*(December), 104167. https://doi.org/10.1016/j.jml.2020.104167

Bogaerts, L., Siegelman, N., Christiansen, M. H., & Frost, R. (2022). Is there such a thing as a 'good statistical learner'?. *Trends in Cognitive Sciences*, *26*(1), 25–37. https://doi.org/10.1016/j.tics.2021.10.012

Burani, C., Marcolini, S., De Luca, M., & Zoccolotti, P. (2008). Morpheme-based reading aloud: Evidence from dyslexic and skilled Italian readers. *Cognition*, *108*(1), 243–262. https://doi.org/10.1016/j.cognition.2007.12.010

Burani, C., Marcolini, S., & Stella, G. (2002). How early does morpholexical reading develop in readers of a shallow orthography? *Brain and Language*, *81*(1-3), 568–586. https://doi.org/10.1006/brln.2001.2548

Chetail, F. (2017). What do we do with what we learn? Statistical learning of orthographic regularities impacts written word processing. *Cognition*, *163*, 103–120. https://doi.org/10.1016/j.cognition.2017.02.015

Clarke, E., & Sherrill-Mix, S. (2017). *ggbeeswarm: Categorical scatter (violin point) plots*. R package version 0.6.0. https://CRAN.R-project.org/package=ggbeeswarm

Coltheart, M., Rastle, K., Perry, C., Langdon, R., & Ziegler, J. C. (2001). DRC: A dual route cascaded model of visual word recognition and reading aloud. *Psychological Review*, *108*(1), 204–256. doi:10.1037/0033-295X.108.1.204

Crepaldi, D., Amenta, S., Mandera, P., Keuleers, E., & Brysbaert, M. (2016). *Frequency estimates from different registers explain different aspects of visual word recognition*. International meeting of the psychonomic society, Granada, 5–8 May.

Crepaldi, D., Rastle, K., Coltheart, M., & Nickels, L. (2010a). 'Fell' primes 'fall', but does 'bell' prime 'ball'? Masked priming with irregularly-inflected primes. *Journal of Memory and Language*, *63*(1), 83–99. https://doi.org/10.1016/j.jml.2010.03.002

Crepaldi, D., Rastle, K., & Davis, C. J. (2010b). Morphemes in their place: Evidence for position-specific identification of suffixes. *Memory & Cognition*, *38*(3), 312–321. https://doi.org/10.3758/MC.38.3.312

Davis, M. H., & Gaskell, G. M. (2009). A complementary systems account of word learning: Neural and behavioural evidence. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *364*(1536), 3773–3800. https://doi.org/10.1098/rstb.2009.0111

Dawson, N., Rastle, K., & Ricketts, J. (2021). Bridging form and meaning: Support from derivational suffixes in word learning. *Journal of Research in Reading*, *44*(1), 27–50. https://doi.org/10.1111/1467-9817.12338

DeBruine, L. M., & Barr, D. J. (2021). Understanding mixed-effects models through data simulation. *Advances in Methods and Practices in Psychological Science*, *4*(1), 1–15. https://doi.org/10.1177/2515245920965119

de Leeuw, J. R., Gilbert, R. A., & Luchterhandt, B. (2023). jsPsych: Enabling an open-source collaborative ecosystem of behavioral experiments. *Journal of Open Source Software*, *8*(85), 5351. https://joss.theoj.org/papers/10.21105/joss.05351

Dowle, M., & Srinivasan, A. (2023). *data.table: Extension of data.frame*. R package version 1.14.8. https://CRAN.R-project.org/package=data.table

Duñabeitia, J. A., Perea, M., & Carreiras, M. (2008). Does*darkness*lead to*happiness*? Masked suffix priming effects. *Language and Cognitive Processes*, *23*(7-8), 1002–1020. https://doi.org/10.1080/01690960802164242

Feldman, L. B., Connor, P. A., & Del Prado Martín, F. (2009). Early morphological processing is morphosemantic and not simply morpho-orthographic: A violation of form-then-meaning accounts of word recognition. *Psychonomic Bulletin andamp*, *16*(4), 684–691. https://doi.org/10.3758/PBR.16.4.684

Feldman, L. B., Kostić, A., Gvozdenović, V., O'connor, P. A., & Del Prado Martín, F. (2012). Semantic similarity influences early morphological priming in Serbian: A challenge to form-then-meaning accounts of word recognition. *Psychonomic Bulletin & Review*, *19*(4), 668–676. https://doi.org/10.3758/s13423-012-0250-x

Fiser, J., & Aslin, R. N. (2001). Unsupervised statistical learning of higher-order spatial structures from visual scenes. *Psychological Science*, *12*(6), 499–504. https://doi.org/10.1111/1467-9280.00392

Fox, J. (2003). Effect displays in *R* for generalised linear models. *Journal of Statistical Software*, *8*(15), 1–27. https://doi.org/10.18637/jss.v008.i15

Frank, S. L., & Bod, R. (2011). Insensitivity of the human sentence-processing system to hierarchical structure. *Psychological Science*, *22*(6), 829–834. https://doi.org/10.1177/0956797611409589

Frost, R., Armstrong, B. C., & Christiansen, M. H. (2019). Statistical learning research: A critical review and possible new directions. *Psychological Bulletin*, *145*(12), 1128–1153. https://doi.org/10.1037/bul0000210

Gathercole, S. E., & Baddeley, A. D. (1993). *Working memory and language*. Psychology Press.

Ginestet, E., Shadbolt, J., Tucker, R., Bosse, M. L., & Deacon, S. H. (2021). Orthographic learning and transfer of complex words: Insights from eye tracking during Reading and learning tasks. *Journal of Research in Reading*, *44*(1), 51–69. https://doi.org/10.1111/1467-9817.12341

Grainger, J., & Beyersmann, E. (2017). Psychology of learning and motivation. *Progress in Brain Research*, *67*, 285–317. https://doi.org/10.1016/bs.plm.2017.03.009

Grainger, J., & Ziegler, J. C. (2011). A dual-route approach to orthographic processing. *Frontiers in Psychology*, *2*, 1–13. https://doi.org/10.3389/fpsyg.2011.00054

Hale, J. (2001). *A probabilistic earley parser as a psycholinguistic model*. 2nd meeting of the north American chapter of the association for computational linguistics, NAACL 2001. 1–7 June 2001. https://doi.org/10.3115/1073336.1073357

Hasenäcker, J., Beyersmann, E., & Schroeder, S. (2016). Masked morphological priming in German-speaking adults and children: Evidence from response time distributions. *Frontiers in Psychology*, *7*, 1–11. https://doi.org/10.3389/fpsyg.2016.00929

Havas, V., Waris, O., Vaquero, L., Rodríguez-Fornells, A., & Laine, M. (2015). Morphological learning in a novel language: A cross-language comparison. *Quarterly Journal of Experimental Psychology*, *68*(7), 1426–1441. https://doi.org/10.1080/17470218.2014.983531

Henson, R. N., & Gagnepain, P. (2010). Predictive, interactive multiple memory systems. *Hippocampus*, *20*(11), 1315–1326. https://doi.org/10.1002/hipo.20857

Hockett, C. F. (1963). The problem of universals in language. In J. H. Greenberg (Ed.), *Universals of language* (pp. 1–29). MIT Press.

Imai, M., & Kita, S. (2014). The sound symbolism bootstrapping hypothesis for language acquisition and language evolution. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *369*(1651), 20130298. doi:10.1098/rstb.2013.0298

Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). Lmertest package: Tests in linear mixed effects models. *Journal of Statistical Software*, *82*(13), 1–26. https://doi.org/10.18637/jss.v082.i13

Lelonkiewicz, J. R., Ktori, M., & Crepaldi, D. (2020). Morphemes as letter chunks: Discovering affixes through visual regularities. *Journal of Memory and Language*, *115*, 104152. https://doi.org/10.1016/j.jml.2020.104152

Lelonkiewicz, J. R., Ktori, M., & Crepaldi, D. (2023). Morphemes as letter chunks: Linguistic information enhances the learning of visual regularities. *Journal of Memory and Language*, *130*, 104411. https://doi.org/10.1016/j.jml.2023.104411

Levy, R. (2008). Expectation-based syntactic comprehension. *Cognition*, *106*(3), 1126–1177. https://doi.org/10.1016/j.cognition.2007.05.006

Longtin, C. M., Segui, J., & Hallé, P. A. (2003). Morphological priming without morphological relationship. *Language and Cognitive Processes*, *18*(3), 313–334. https://doi.org/10.1080/01690960244000036

Mak, M. H. C., Hsiao, Y., & Nation, K. (2021). Anchoring and contextual variation in the early stages of incidental word learning during reading. *Journal of Memory and Language*, *118*, 104203. https://doi.org/10.1016/j.jml.2020.104203

Marelli, M., Amenta, S., & Crepaldi, D. (2015). Semantic transparency in free stems: The effect of orthography-semantics consistency on word recognition. *Quarterly Journal of Experimental Psychology*, *68*(8), 1571–1583. https://doi.org/10.1080/17470218.2014.959709

Marelli, M., & Baroni, M. (2015). Affixation in semantic space: Modeling morpheme meanings with compositional distributional semantics. *Psychological Review*, *122*(3), 485–515. https://doi.org/10.1037/a0039267

McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: I. An account of basic findings. *Psychological Review*, *88*(5), 375–407. doi:10.1037/0033-295X.88.5.375

Merkx, M., Rastle, K., & Davis, M. H. (2011). The acquisition of morphological knowledge investigated through artificial language learning. *Quarterly Journal of Experimental Psychology*, *64*(6), 1200–1220. https://doi.org/10.1080/17470218.2010.538211

Nation, I. (2006). How large a vocabulary is needed for reading and listening? *The Canadian Modern Language Review*, *63*(1), 59–82. https://doi.org/10.3138/cmlr.63.1.59

Nation, K. (2008). Learning to read words. *Quarterly Journal of Experimental Psychology*, *61*(8), 1121–1133. https://doi.org/10.1080/17470210802034603

New, B., Brysbaert, M., Segui, J., Ferrand, L., & Rastle, K. (2004). The processing of singular and plural nouns in French and English. *Journal of Memory and Language*, *51*(4), 568–585. https://doi.org/10.1016/j.jml.2004.06.010

Ouellette, G. P. (2006). What's meaning got to do with it: The role of vocabulary in word reading and reading comprehension. *Journal of Educational Psychology*, *98*(3), 554–566. https://doi.org/10.1037/0022-0663.98.3.554

Pedersen, T. L., Shemanarev, M., Juricic, T., Marusinec, M., Garrett, S., & Posit Software, P. B. C. (2023). ragg: Graphic devices based on AGG (R package version 1.2.5). https://CRAN.R-project.org/package=ragg

Plaut, D. C., McClelland, J. L., Seidenberg, M. S., & Patterson, K. (1996). Understanding normal and impaired word Reading: Computational principles in quasi-regular domains. *Psychological Review*, *103*(1), 56–115. https://doi.org/10.1037/0033-295x.103.1.56

Ram, K., Wickham, H., Richards, C., & Baggett, A. (2019). wesanderson: A Wes Anderson palette generator (R package version 0.3.6). https://CRAN.R-project.org/package=wesanderson

Ramscar, M., Hendrix, P., Love, B., & Baayen, R. H. (2013). Phonological and Phonetic considerations of Lexical Processing. *The Mental Lexicon*, *8*(3), 450–481. https://doi.org/10.1075/ml.8.3.08ram

Rastle, K., Davis, M. H., & New, B. (2004). The broth in my brother's brothel: Morpho-orthographic segmentation in visual word recognition. *Psychonomic Bulletin & Review*, *11*(6), 1090–1098. https://doi.org/10.3758/bf03196742

R Core Team. (2025). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. https://www.R-project.org/

Rescorla, R. A. (1988). Pavlovian conditioning: It's not what you think it is. *American Psychologist*, *43*(3), 151–160. doi:10.1037/0003-066X.43.3.151

Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, *274*(5294), 1926–1928. https://doi.org/10.1126/science.274.5294.1926

Sapir, E. (1929). A study in phonetic symbolism. *Journal of Experimental Psychology*, *12*(3), 225–239. https://doi.org/10.1037/h0070931

Schmalz, X., Moll, K., Mulatti, C., & Schulte-Körne, G. (2019). Is statistical learning ability related to reading ability, and if so, why? *Scientific Studies of Reading*, *23*(1), 64–76. https://doi.org/10.1080/10888438.2018.1482304

Seidenberg, M. S. (1987). Sublexical structures in visual word reccognition: Access units or orthographic redundancy? In M. Coltheart (Ed.), *Attention & performance XII: The psychology of Reading* (pp. 245–263). Erlbaum.

Shanks, D. R. (1995). Is human learning rational? *The Quarterly Journal of Experimental Psychology Section A*, *48*(2), 257–279. https://doi.org/10.1080/14640749508401390

Siegelman, N. (2020). Statistical learning abilities and their relation to language. *Language and Linguistics Compass*, *14*(3), 1–19. https://doi.org/10.1111/lnc3.12365

Siegelman, N., Bogaerts, L., & Frost, R. (2017). Measuring individual differences in statistical learning: Current pitfalls and possible solutions. *Behavior Research Methods*, *49*(2), 418–432. https://doi.org/10.3758/s13428-016-0719-z

Siegelman, N., & Frost, R. (2015). Statistical learning as an individual ability: Theoretical perspectives and empirical evidence. *Journal of Memory and Language*, *81*, 105–120. https://doi.org/10.1016/j.jml.2015.02.001

Siegelman, N., Rueckl, J. G., Steacy, L. M., Frost, S. J., van den Bunt, M., Zevin, J. D., Seidenberg, M. S., Pugh, K. R., Compton, D. L., & Morris, R. D. (2020). Individual differences in learning the regularities between orthography, phonology and semantics predict early reading skills. *Journal of Memory and Language*, *114*, 104145. https://doi.org/10.1016/j.jml.2020.104145

Smith, L. B., Colunga, E., & Yoshida, H. (2010). Knowledge as process: Contextually cued attention and early word learning. *Cognitive Science*, *34*(7), 1287–1314. https://doi.org/10.1111/j.1551-6709.2010.01130.x

Solaja, O., & Crepaldi, D. (2024). The role of morphology in novel word learning: A registered report. *Royal Society Open Science*, *11*, 230094. https://doi.org/10.1098/rsos.230094

Taft, M., & Forster, K. I. (1975). Lexical storage and retrieval of prefixed words. *Journal of Verbal Learning and Verbal Behavior*, *14*(6), 638–647. doi:10.1016/S0022-5371(75)80051-X

Taft, M., & Nguyen-Hoan, M. (2010). A sticky stick? The locus of morphological representation in the lexicon. *Language and Cognitive Processes*, *25*(2), 277–296. https://doi.org/10.1080/01690960903043261

Tamminen, J., Davis, M. H., & Rastle, K. (2015). From specific examples to general knowledge in language learning. *Cognitive Psychology*, *79*, 1–39. https://doi.org/10.1016/j.cogpsych.2015.03.003

Treiman, R., Kessler, B., Boland, K., Clocksin, H., & Chen, Z. (2018). Statistical learning and spelling: Older prephonological spellers produce more wordlike spellings than younger prephonological spellers. *Child Development*, *89*(4), 1–13. https://doi.org/10.1111/cdev.12893

Ullman, M. T. (2001). The declarative/procedural model of lexicon and grammar. *Journal of Psycholinguistic Research*, *30*(1), 37–69. https://doi.org/10.1023/A:1005204207369

Viviani, E., & Crepaldi, D. (2022). Masked morphological priming and sensitivity to the statistical structure of form–to–meaning mapping in L2. *Journal of Cognition*, *5*(1), 30. https://doi.org/10.5334/joc.221

Wagner, A. R., & Rescorla, R. A. (1972). Inhibition in Pavlovian conditioning: Application of a theory. In R. A. Boakes & M. S. Halliday (Eds.), *Classical conditioning II: Current research and theory* (pp. 301–336). Appleton-Century-Crofts.

Wei, T., Simko, V., Levy, M., Xie, Y., Jin, Y., Zemla, J., Freidank, M., Cai, J., & Protivinsky, T. (2021). corrplot: Visualization of a correlation matrix (R package version 0.92). https://CRAN.R-project.org/package=corrplot

Wickham, H. (2016). *Ggplot2: Elegant graphics for data analysis*. Springer-Verlag New York. https://ggplot2.tidyverse.org

Wickham, H., Averick, M., Bryan, J., Chang, W., D'Agostino McGowan, L., François, R., Grolemund, G., Hayes, A., Henry, L., Hester, J., Kuhn, M., Pedersen, T. L., Miller, E., Bache, S. M., Müller, K., Ooms, J., Robinson, D., Seidel, D. P., Spinu, V., … Yutani, H. (2019). Welcome to the tidyverse. *Journal of Open Source Software*, *4*(43), 1686. https://doi.org/10.21105/joss.01686

Wilke, C. O. (2020). *cowplot: Streamlined plot theme and plot annotations for 'ggplot2'*. R package version 1.1.1. https://CRAN.R-project.org/package=cowplot

Yablonski, M., & Ben-Shachar, M. (2016). Linguistic perspectives on morphological processing. *The Mental Lexicon*, *11*, 277–307. https://doi.org/10.1075/ml.11.2.05yab

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. https://doi.org/10.3758/PBR.15.5.971