Exploring Semantic Relatedness Judgments in the Structure of a Semantic Network[1]

Polyphony J. Bruna (pbruna@vassar.edu)
Department of Cognitive Science, Vassar College
Poughkeepsie, NY 12604 USA

Abstract

Drawing upon work by De Deyne et al. (2016), I explore a model of spreading activation through a semantic network in regards to how different kinds of semantic relationships are encoded in said network. In particular, I examine the contribution of indirect pathways through the network to explain differences in similarity judgments of sensorimotor and linguistic relationships between pairs of words. I propose that the structure of a semantic network encodes properties that distinguish these two types of semantic relationships that are not revealed by measures of association strength that only examine direct connections within the network. A cosine similarity measure extracted from a spreading activation model is compared to a measure of association strength in accounting for observed similarity judgments, and a model for examining the differential contributions of various random walk pathways through a semantic network in the encoding of sensorimotor and linguistic semantic relationships is presented.

Keywords: similarity; relatedness; semantic network; spreading activation; random walk; linguistic; sensorimotor; modeling

Introduction

Semantic networks encode relationships between words as weighted connections between nodes that represent concepts. The properties of such networks can be explored to examine the structure of semantic knowledge. De Deyne et al. (2016) develop a semantic network from word association data and use a spreading activation model to explore weak semantic associations. It is unclear on precisely what basis people make judgments about similarity, and this is especially the case when it comes to items which are themselves highly dissimilar. This leads to the possibility that such ratings are highly inconsistent and uninterpretable. De Deyne et al. challenge this assumption and demonstrate that structure underlying such judgments may be recovered in the indirect connections of spreading activation through a semantic network.

The authors first present behavioral evidence from a series of experiments which indicate that systematic patterns can be found in peoples' similarity judgments between highly dissimilar items. They ask participants to choose the most similar pair out of three items and find significant preferences despite all three items being highly dissimilar (p. 12). They then develop a model in order to offer an explanation for how these patterns may be reflected in the structure of a semantic network. Using free association data from N = 12,428 cue words with up to three free responses given for each cue word (De Deyne & Storms, 2008; De Deyne et al., 2011), they create both a sparsely and a densely connected semantic network, which make use of only the first free response and all three free responses given by a participant to a cue, respectively. This network is represented by an adjacency matrix of the associative strengths of word pairs, P. The matrix is transformed into a random walk graph, G_{rw}, to represent spreading activation through the network. This transformation corresponds to the culmination of many random walks through the network in which the activation of a node as it is "passed through" by a walk initiates more random walks. G_{rw} represents the number and length of the paths through the network that connect any two nodes (De Deyne et al., 2016, p. 15). The random walk is controlled by the parameter, α, indicating the rate of decay of activation (p. 16, 19).[2] The more similar the distribution of paths through the network is for any two words, the more similar those words are considered to be (p. 15). This measure is calculated from G_{rw} as the cosine similarity between any two words (this is described in more depth below).

Cosine similarity calculated from G_{rw} was found to predict the judgments of each of their behavioral studies better than cosine similarities calculated from only the distributional overlap of connections in the semantic network without spreading activation, with the dense network outperforming the sparse network. Similarities incorporating indirect paths using the spreading activation model performed just as well as similarities that did not include indirect paths only for the dense network on responses to highly similar items (p. 18). The authors conclude that the underlying semantic structures that may be responsible for producing systematic similarity

[1] This work was completed as a final project in a cognitive science modeling course and subsequent independent study project. Modeling techniques were employed to expand upon and examine results obtained from a prior behavioral study (Bruna, 2020) completed as a final project for a cognitive science seminar course.

[2] $G_{rw} = \sum_{r=0}^{\infty}(\alpha P)^r = (I - \alpha P^{-1})$

judgments on highly dissimilar items may be represented by the indirect connections traced out by spreading activation through a semantic network (p. 19).

The authors follow up on these results by creating a version of the spreading activation model in which each possible path through the network is assigned a parameter that differentially weights the influence of the different path types (p. 21). Training these parameters on the behavioral data from their experiments revealed a significant contribution of longer paths in the experiments examining similarity judgments on highly dissimilar items and a significant contribution of direct paths alongside indirect paths in the experiment examining judgments on similar items (p. 22). From this the authors are able to conclude that the ability of their spreading activation model to reflect the patterns found in the similarity judgments of dissimilar items may be attributed to the influence of indirect pathways through the semantic network which are not revealed when only direct pathways are considered (p. 22). De Deyne et al. conclude that similarity judgments encompass not merely lexical association but also underlying semantic structures (p. 23).

To expand upon this finding, I explore the indirect semantic pathways traced out by De Deyne et al.'s (2016) spreading activation model in order to explore the representation of different kinds of semantic knowledge on their account. In particular, I compare the representation of associations that arise from sensorimotor knowledge, the kind of relational knowledge that De Deyne et al. (2016) emphasize in their exploration of weak similarities, to linguistic associations produced purely by lexical co-occurrence.

The relational structure of language has been proposed as a possible source of semantic knowledge alongside embodied experiences. Evidence that distributional models are able to extract meaning from the statistical patterns present in language, also called natural language statistics (NLS), supports the idea that language itself can contribute to the meaning of our concepts (Lupyan & Lewis, 2019, p. 1324). These models extract semantic similarities between words based on the frequencies with which they appear in similar contexts. For example, a machine algorithm called word2vec extracts analogical reasoning from these distributional statistics (e.g. applying the relation ANIMAL → WOLF to FISH yields SHARK; however, applying the relation ANIMAL→ DOG yields GOLDFISH (p. 1325)). Furthermore, Lewis et al. (2019) found that judgments of animal similarity on the basis of shape, texture, and color made by distributional algorithms significantly correlated with those made by blind participants on an analogous task.

An upshot of the proposal that the structure of language contributes to semantic knowledge is that it provides one possible source of abstract knowledge, which cannot readily be explained through embodied representations. Abstract words are those that cannot be described by merely pointing out a referent or enacting its meaning. Lupyan and Winter (2018) emphasize the challenge to account for abstract meaning by asserting that abstract words make up most of our communication. They observe that in statements only five words long there is a 95% chance of encountering a word that is as abstract as "freedom" (p. 3). This highlights the implausibility that semantics can be entirely reduced to sensorimotor, affective, and situational experiences; however, amodal theories that merely posit that our semantics map onto innate concepts fail to adequately address how these concepts arise (p. 4). Attributing abstract knowledge at least in part to linguistic structure provides one possible response to this problem.

Furthermore, Marques and Nunes (2012) propose that sensorimotor and linguistic relational knowledge may differentially contribute to constructing abstract and concrete concepts. They observed that the strongest word associates for abstract words tend to be related through linguistic knowledge, whereas the strongest word associates for concrete words do not significantly differ between being related through sensorimotor or linguistic knowledge (Marques & Nunes, 2012, p. 1270–1271). These findings support the conclusion that language occupies an influential role in the acquisition and representation of abstract knowledge. Moreover, it suggests that the potential differential representation of sensorimotor relational knowledge compared to linguistic relational knowledge in a spreading activation network model of semantic knowledge, such as that developed by De Deyne et al. (2016), may have bearing on the encoding and retrieval of abstract compared to concrete concepts.

Behavioral evidence from Bruna (2020) suggests that a deviance between the representation of sensorimotor and linguistic relational knowledge is revealed in judgments of similarity that are not reflected in an association strength measure. In this study, participants were asked to judge the strength of the relation between pairs of words using a 13-point slider marked only with a "-" symbol on the left-hand side and a "+" symbol on the right-hand side (p. 7). Word pairs were classified as either sensorimotor associates or linguistic associates. Words were presented simultaneously as clip-art style images to increase recognizability and iconicity, and all words were chosen so that they could be unambiguously represented by an image (p. 9). Participants completed this task while simultaneously completing either a linguistic interference task, a spatial interference task, or no interference task (p. 6).

The purpose of this study was to explore potential differential access to different types of relational knowledge across interference conditions. To compare these differences across interference conditions, the sensorimotor and linguistic word pairs were controlled on association strength measured as response probability (p. 9). This was expected to produce no significant difference in the relational strength ratings between sensorimotor and linguistic pairs in the control condition. The control was unsuccessful: in the control condition, word pairs identified as sensorimotor associates produced significantly higher strength ratings compared to linguistic associates (p. 11). It was hypothesized that these results may be attributable to the fact that pairs were presented as images despite the fact that the clip-art style of the images was chosen to evoke a general representation of the intended concept that would be comparable to the kind of representation evoked by a word. This would indicate that variation in the domain of input produces differential access semantic knowledge (p. 16). The following model explores the different hypothesis that the structure of a spreading activation model of a semantic network, such as that proposed by De Deyne et al. (2016), encodes the semantic knowledge that is responsible for the differential rating of sensorimotor and linguistic pairs, and that this structure is not revealed by the association strength used as a control measure in Bruna (2020).

The Model

Computations were performed in the R environment (v. 3.5) (R Core Team, 2020), using the ggplot2 (Wickham, 2016), dplyr (Wickham et al., 2020), tidyverse (Wickham et al., 2019), readr (Wickham, et al., 2018), BayesFactor (Morey & Rouder, 2018), Matrix (Bates & Maechler, 2017), tictoc (Izrailev, 2014), igraph (Csardi & Nepusz, 2006), and optimParallel (Gerber & Furrer, 2019) packages. Scripts and other supplemental materials may be found at: https://osf.io/3e2tq/. Scripts for the original model developed by De Deyne et al. (2016) may be found at: https://github.com/SimonDeDeyne/SWOWEN-2018.
Finally, the free association data for generating the semantic network used was provided by De Deyne et al. (2018) and may be found at: https://smallworldofwords.org/en/project/research.

Free association data using N = 12,292 cue words with up to three free responses given for each cue word (De Deyne et al., 2018) is used to generate an adjacency matrix that represents a network in which words are represented as nodes with weighted connections between them. Only the first free response given to a cue was used to construct the sparsely connected network described by De Deyne et al. (2016) above. All three free responses given to a cue were used to construct the densely connected network. Although the densely connected network was shown to marginally outperform the sparsely connected network in De Deyne et al., both the sparsely connected network and the densely connected network are compared in the present work. This matrix is constrained to include only cues that were also given as free association responses at least once, thus ensuring both in-going and out-going connections for every node in the network. The connection weights between words are transformed using the positive pointwise mutual information (PMI^+) measure,[3] which ensures that responses that are very frequently given for many cues are considered less informative than responses that are given frequently for only select cues (De Deyne et al., 2016, p. 14). The similarity between words is assessed as the distributional overlap of shared paths through the network, including indirect neighbors up to a path length of three (a path of length one describes a direct connection between two nodes in the network; a path of length two describes a path between two nodes that is mediated by a third node; a path of length three describes a path between two nodes that is mediated by a third and a fourth node) (p. 21). A random walk procedure is first implemented to represent spreading activation through the network, as explained above. This is achieved by transforming the adjacency matrix, P, into the random walk graph, G_{rw}, using the following equation (p. 21):

$$
\begin{aligned}
G_{rw} &= \beta_1 P + && \text{(paths of length 1)} \\
&\beta_2 PP^T + \beta_3 P^2 + \beta_4 P^T P + && \text{(paths of length 2)} \\
&\beta_5 P^3 + \beta_6 P^2 P^T + \beta_7 PP^T P + \beta_8 P^T P^2 && \text{(paths of length 3)}
\end{aligned}
$$

In this equation, each possible path through the network is assigned a parameter, β, which weights the influence of the path on the spreading activation through the network. The parameter values fall between zero and one and are constrained such that the sum of all eight parameters is equal to one. As described above, G_{rw} represents the number and length of the paths through the network that connects any two nodes. Finally, the similarity between any two words is calculated using the cosine measure of similarity. This measure describes the similarity between two vectors as the cosine of the angle between them. Two vectors which are identical in orientation will produce an angle of 0° and hence a cosine similarity of cos(0) = 1. To calculate the cosine similarity between any two words, the rows of the two words are extracted from G_{rw} and the dot product of these vectors is calculated after the vectors have been normalized using the L2 norm. Normalizing the vectors accounts for differential frequency across the two vectors (p. 17). Because the random walk procedure introduces indirect activation paths through the network, calculating cosine similarity from G_{rw} is able to represent the influence of underlying semantic structures in the network that are not revealed by associative strength measures which only take into account directly neighboring nodes. This model will allow for the exploration of how word pairs that are associated through sensorimotor relational knowledge compared to word pairs that are associated through linguistic relational knowledge may be differentially represented in the semantic network via the different path types.

Fitting the Model

The word pairs of interest were those used in Bruna (2020). These were 16 pairs, divided into eight sensorimotor associations and eight linguistic associations. Sensorimotor associations were situationally related (e.g. BIRD → TREE) and linguistic associations consisted of forward and backward associates (e.g. STRAWBERRY → BLONDE).

[3] $PMI^+(p_{i|j}) = \max\left(0, log_2\left(\frac{p(i|j)}{\sum_j^n p(i|j)}\right)\right)$

The pairs were exclusively situationally or linguistically related so that no linguistically related associate that was also non-linguistically related (e.g. WINE → BARREL), or vice versa, was chosen (Bruna, 2020, p. 9).

These associations were generated from eight cue words selected from the 47 concrete words used by Marques and Nunes (2012). These words were rated as highly concrete (ratings > 6.45; Paivio et al, 1968) and were controlled for word frequency, familiarity, length (from Clark & Paivio, 2004), and orthographic neighbors (http://neighborhoodsearch.wustl.edu/Neighborhood/NeighborHome.asp) (Marques & Nunes, 2012, p. 1268). Two associates were chosen for each word from the University of South Florida free association norms database (Nelson et al., 2004). One of these associates was related to the target word by a sensorimotor relation and the other was related by a linguistic relation. For each cue word, the two associates were controlled on the strength of association, measured as the proportion of individuals who produced the associate when prompted with the target in a free association task; however, association strength differed between target words and ranged from 0.01 to 0.177 (Bruna, 2020, p. 9).

Despite controlling for association strength, contrary to expectations the sensorimotor and linguistic associations produced significantly different relational strength ratings in the control condition of the experiment (p. 11). This pattern was confirmed in the semantic network of De Deyne et al. (2016). With the exception of one cue word, the association strength for each pair calculated using the free association data from De Deyne et al. (2018) similarly showed support for there being no difference in the sensorimotor associate compared to the linguistic associate for a given cue word (excluding the outlier, $BF_{01} = 3.102437$; including the outlier, $BF_{01} = 2.366571$), and association strengths ranged from 0.00 to 0.26. In addition to this, with the exception of one word cue (different from the aforementioned), cosine similarities calculated for each pair from the version of De Deyne et al.'s (2016) random walk model that makes use of a single decay parameter rather than parameters weighting each possible path through the model demonstrate support for there being higher similarities for the sensorimotor associate compared to the linguistic associate for each cue (excluding the outlier, $BF_{10} = 6.372362$; including the outlier, $BF_{10} = 5.10439$). This indicates that the spreading activation model captures semantic knowledge regarding sensorimotor and linguistic relations that is not revealed in association strength alone. Furthermore, the higher relational strength ratings for sensorimotor associations found in the control condition of Bruna (2020) may be attributable to the structure of De Deyne et al.'s (2016) semantic network rather than the shift from linguistic to pictorial domains of input, as initially proposed in Bruna (2020).

In order to explore this possibility, the spreading activation model weighting each of the eight possible path types described above is fit to the response data collected by Bruna (2020). Parallel Nelder-Mead optimization is used to estimate each of the eight path type parameters of the model using random starting values for the sensorimotor word pairs and the linguistic word pairs separately. Out of the eight word cues used in Bruna (2020), the aforementioned two exception word cues that did not conform to the pattern of interest were excluded prior to model training, resulting in a total of six pairs in the sensorimotor set and six pairs in the linguistic set. The cosine similarity ratings calculated from the model were compared with the average rating from the behavioral data Bruna (2020) for each word pair to produce a correlation coefficient that was used as a measure of discrepancy during model fitting.

Results

Model fitting was performed using Hopper, a high-performance computing cluster hosted by Vassar College. Parameter values are given in Table 1.

The model reached convergence in all runs and produced a correlation coefficient between the cosine similarity of each word pair in the model and the observed behavioral judgments of these word pairs ranging from 0.75 to 0.78. The distribution of weights across the eight parameters do not appear to significantly differ across the sparsely and densely connected networks. Furthermore, the parameters which produce cosine similarities between word pairs in the model that match the pattern of observed judgments of semantic relatedness do not appear to differ significantly across sensorimotor word pairs and linguistic word pairs. In other words, approximately the same set of parameter weights produces measures of cosine similarity that accurately describe the observed pattern of semantic relatedness judgments in both the linguistic word pairs and the sensorimotor word pairs. Furthermore, this set of parameter weights supports both direct and indirect paths between words in both cases.

These results do not support the hypothesis that linguistically associated word pairs should rely more heavily upon direct paths compared to word pairs that are associated through sensorimotor knowledge. This hypothesis was supported by the notion that words that are only linguistically associated should derive their association from lexical co-occurrence alone. To use an example from De Deyne et al. (2016), a consistent judgment of similarity between the words ATHLETE and BREATH was found in their behavioral studies of weak similarity. These words are not direct associates, but as the authors point out it is not difficult to produce a chain of associative reasoning connecting the two: athletes physically exert themselves, which causes them to breathe heavily (p. 14–15). This kind of reasoning is possible because of the sensorimotor knowledge that connects the two concepts. However, this reasoning should not be possible for word pairs that are exclusively related through lexical co-occurrence. For example, BIRD produces BATH in a free association task, but it is difficult to imagine this connection being mediated by other words. Thus, one might expect that sensorimotor associates may be connected through direct and indirect paths in the semantic network but

Table 1 The optimized parameter values for the eight paths through the network using the densely connected network. A full table including parameter values from the sparsely connected network may be found at: https://osf.io/3e2tq/.

	Sensorimotor	Linguistic
β_1	1.563226e-01	2.597393e-01
β_2	5.196752e-02	6.925390e-03
β_3	9.937230e-05	3.896904e-07
β_4	7.344489e-01	7.146740e-02
β_5	1.763123e-03	1.553563e-06
β_6	2.871386e-06	3.406172e-01
β_7	5.537012e-02	3.677076e-08
β_8	2.552427e-05	3.212488e-01

that linguistic associates should be connected dominantly by direct paths. This also reveals that one might expect directionality to play an important role in how linguistic associates are related within a semantic network. My results indicate that the differences in observed semantic relatedness judgments across words that are related through sensorimotor association and words that are related through linguistic association is not explained by differences in the structural features of a semantic network. This consideration is taken up in the discussion section below.

Discussion

A spreading activation model illuminates structural properties of a semantic network that are not apparent when merely considering direct semantic neighbors or association strength measures. That people tend to respond in predictable ways to weak similarity judgments, judgments which at first glance would not appear to involve any systematic basis for response, demands an exploration of what underlying structure might be responsible for these patterns. De Deyne et al. (2016) demonstrate that similarity measures that incorporate the various indirect pathways through a semantic network, as revealed by a spreading activation model, capture the response patterns made by people.

These findings implicate that different kinds of relational knowledge may be encoded in the various possible paths through a semantic network. For example, it would be expected that sensorimotor relationships such as that between BIRD and TREE may arise not only from a direct connection between these two concepts but also be supported by indirect connections that are mediated by similarly related notions. However, it would not be expected for this to be the case with word associations that arise solely from linguistic structure via lexical co-occurrence. For example, BIRD may elicit BATH in a free response task, but it seems unlikely to expect that this relationship would be mediated by other notions. While the associative chain BIRD → NEST → TREE appears intuitively plausible, BIRD → ? → BATH does not.

It was found that two of the word cues examined in this study did not exhibit the pattern of interest wherein sensorimotor and linguistic associations that were matched on association strength revealed a stronger similarity score for the sensorimotor relationship when measured as cosine similarity on a spreading activation model. Excluding these pairs allows for the exploration of this pattern; however, what they reveal is that the results of this study cannot be generalized to all sensorimotor and linguistic relations. Further steps include exploring these variations in the different types of semantic relationships within a more expansive set of stimulus pairs.

By modeling spreading activation through a semantic network, semantic relations may be examined in the context of their situatedness within larger semantic structures. This provides a richer understanding compared to the examination of these relationships on the basis of association strength measured as response probability alone. This is made apparent in the use of similarity measures derived from a spreading activation model in De Deyne et al. (2016) to uncover systematic relationships in weak similarity judgments.

Evidence from Bruna (2020) suggests that association strengths do not capture all of the information about the relationship between pairs of words necessary to make predictions regarding participant judgments of relatedness of these pairs. The present study reveals that the difference in observed semantic relatedness judgments on sensorimotor and linguistic word pairs is also not explained by differences in the structural properties of a semantic network as revealed by activation of various direct and indirect paths through the network. Importantly, counter to the hypothesis, results would seem to suggest that word pairs that are related through linguistic knowledge (in this case, lexical co-occurrence) do not privilege a direct path connection in the model. Instead, in both the case of linguistic word pairs and sensorimotor word pairs, both direct and indirect pathways support word activation. It is possible that these results may support a sort of linguistic shortcut theory wherein verbal labels, rather than replacing, act as a shorthand for richer sensorimotor information (Connell, 2019). Furthermore, such an account treats the act of using language itself as a sensorimotor process that contributes to semantic knowledge. Such a claim blurs any sharp distinction between linguistic knowledge and sensorimotor knowledge, a conclusion that appears to be supported by the present results.

Although such an explanation was not found here, development of a computational explanation for the behavioral differences observed in Bruna (2020) would allow for the possibility of a computational reformulation of the notion of "different types of semantic relationships" as it has been used as a manipulation in studies such as Marques and

Nunes (2012) and Bruna (2020) and would present a basis on which to ground such manipulations in future studies.

References

Bates, D., & Maechler, M. (2017). Matrix: Sparse and Dense Matrix Classes and Methods. R package version 1.2-12. https://CRAN.R-project.org/package=Matrix

Bruna, P. J. (2020). *Exploring the Role of Language in Relational Concept Formation* [Unpublished manuscript]. Department of Cognitive Science, Vassar College.

Clark, J. M. & Paivio, A. (2004). Extensions of the Paivio, Yuille, and Madigan (1968) norms. *Behavior Research Methods, Instruments, & Computers, 36*, 371–383. doi:10.3758/BF03195584

Connell, L. (2019). What have labels ever done for us? The linguistic shortcut in conceptual processing. *Language, Cognition and Neuroscience, 34*(10), 1308–1318. doi: 10.1080/23273798.2018.1471512

Csardi, G., & Nepusz, T. (2006). The igraph software package for complex network research, *InterJournal*, *Complex Systems*. https://igraph.org

De Deyne, S., Navarro, D. J., Perfors, A., Brysbaert, M., & Storms, G. (2018). The "Small World of Words" English word association norms for over 12,000 cue words. *Behavior Research Methods.* DOI 10.3758/s13428-018-1115-7.

De Deyne, S., Navarro D. J., Perfors, A., & Storms, G. (2016). Structure at every scale: A semantic network account of the similarities between very unrelated concepts. *Journal of Experimental Psychology: General, 145*, 1228–1254. https://doi.org/10.1037/xge0000192.

De Deyne, S., & Storms, G. (2008b). Word associations: Norms for 1,424 Dutch words in a continuous task. *Behavior Research Methods, 40*, 198–205.

De Deyne, S., Voorspoels,W., Verheyen, S., Navarro, D. J., & Storms, G. (2011). Graded structure in adjective categories. In L. Carlson, C. Hölscher, & T. F. Shipley (Eds.), *Proceedings of the 33rd Annual Conference of the Cognitive Science Society* (pp. 1834–1839). Austin, TX: Cognitive Science Society.

Gerber, F., & Furrer, R. (2019). optimParallel: An R package providing a parallel version of the L-BFGS-B optimization method. *The R Journal, 11*(1), 352–358. https://doi.org/10.32614/RJ-2019-030

Izrailev, S. (2014). tictoc: Functions for timing R scripts, as well as implementations of Stack and List structures. R package version 1.0. https://CRAN.R-project.org/package=tictoc

Lewis, M., Zettersten, M., & Lupyan, G. (2019). Distributional semantics as a source of visual knowledge. *Proceedings of the National Academy of Sciences. 116*. 201910148.10.1073/pnas.1910148116.

Lupyan, G., & Lewis, M. (2019). From words-as-mappings to words-ascues: the role of language in semantic knowledge, *Language, Cognition and Neuroscience, 34*(10), 1319–1337, https://doi.org/10.1080/23273798.2017.1404114

Lupyan, G., & Winter, B. (2018). Language is more abstract than you think, or, why aren't languages more iconic? *Philosophical Transactions of the Royal Society B Biological Sciences, 373*(1752). https://doi.org/10.1098/rstb.2017.0137

Marques, J. F., & Nunes, L. D. (2012). The contributions of language and experience to the representation of abstract and concrete words: Different weights but similar organizations. *Memory and Cognition, 40,* 1266–1275. https://doi.org/10.3758/s13421-012-0220-6

Morey, R. D., & Rouder, J. N. (2018). BayesFactor: Computation of Bayes Factors for Common Designs. R package version 0.9.12-4.2. https://CRAN.R-project.org/package=BayesFactor

Mullen, K., Ardia, D., Gil, D., Windover, D., & Cline, J. (2011). DEoptim: An R Package for Global Optimization by Differential Evolution. *Journal of Statistical Software, 40*(6), 1–26. http://www.jstatsoft.org/v40/i06/.

Nelson, D. L., McEvoy, C. L., & Schreiber, T. A. (2004). The University of South Florida free association, rhyme, and word fragment norms. *Behavior Research Methods, Instruments, & Computers, 36,* 402–407. doi:10.3758/BF03195588

Paivio, A., Yuille, J. C., & Madigan, S. A. (1968). Concreteness, imagery and meaningfulness values for 925 nouns. *Journal of Experimental Psychology, 76*(1, Pt. 2), 1–25. doi:10.1037/h0025327

R Core Team (2020). R: A language and environment for statistical computing. R Foundation for Statistical Computing. Vienna, Austria. URL https://www.R-project.org/

Wickham, H. (2016). ggplot2: Elegant Graphics for Data Analysis. Springer-Verlag: New York.

Wickham, H., François, R., Henry, L., & Müller, K. (2020). dplyr: A Grammar of Data Manipulation. R package version 0.8.4. https://CRAN.R-project.org/package=dplyr

Wickham, H., Hester, J., & Francois, R. (2018). readr: Read Rectangular Text Data. R package version 1.3.1. https://CRAN.R-project.org/package=readr

Wickham et al., (2019). Welcome to the tidyverse. *Journal of Open Source Software, 4*(43). https://doi.org/10.21105/joss.01686