

© 2019, American Psychological Association. This paper is not the copy of record and may not exactly replicate the final, authoritative version of the article. Please do not copy or cite without authors permission. The advanced online publication is available at <https://doi.org/10.1037/xlm0000793>, and the paper will appear in print in the Journal of Experimental Psychology: Learning, Memory, and Cognition. The final article will be available, upon publication, via its DOI.

Distant Connectivity and Multiple-Step Priming in Large-Scale Semantic Networks

Abhilasha A. Kumar¹, David A. Balota¹, Mark Steyvers²

¹ Department of Psychological & Brain Sciences, Washington University in St Louis, MO

² Department of Cognitive Sciences, University of California, Irvine, CA

Authors' Note

All data and analysis scripts used in this project and have been made available at <https://github.com/abhilasha-kumar/Distant-Semantic-Connectivity>. Portions of this work were presented at the 59th Annual Meeting of the Psychonomic Society (2018) in New Orleans, LA, at the 41st Annual Meeting of the Cognitive Science Society in Montreal, Canada, and published in the Cognitive Science Conference Proceedings. Correspondence concerning this article should be addressed to Abhilasha A. Kumar, Department of Psychological and Brain Sciences, Washington University in St. Louis, Campus Box 1125, One Brookings Drive, St. Louis, MO 63130. E-mail: abhilasha.kumar@wustl.edu

Abstract

We examined 3 different network models of representing semantic knowledge (5018-word directed and undirected step distance networks, and an association-correlation network) to predict lexical priming effects. In Experiment 1, participants made semantic relatedness judgments for word pairs with varying path lengths. Response latencies for judgments followed a quadratic relationship with network path lengths, replicating and extending a recent pattern reported by Kenett, Levi, Anaki, and Faust (2017, *Journal of Experimental Psychology: Learning, Memory & Cognition*) for an 800-word association-correlation network in Hebrew. In Experiment 2, participants identified target words in a progressive demasking task, immediately following a briefly presented prime (120 ms). Response latencies to identify the target showed a linear trend for all network path lengths. Importantly, there were statistically significant differences between relatively distant words in the step distance networks, e.g., path lengths 4 and beyond, suggesting that association networks can indeed capture distant functional semantic relationships. Additional comparisons with two distributional models (LSA and *word2vec*) suggested that distributional models also successfully predicted response latencies, although there appear to be fundamental differences in the types of semantic relationships captured by the different models.

Keywords: semantic priming, semantic networks, semantic space models

Understanding language requires the retrieval of meaning from underlying semantic representations of words. A standard model of retrieving meaning from semantic memory involves a spread of activation, such that activation spreads from one concept to related concepts along associative/semantic pathways (e.g., Collins & Loftus, 1975; Collins & Quillian, 1969). A common finding that directly follows from the spreading-activation account is that processing a particular word (e.g., cat) facilitates processing of a related word (e.g., dog), a phenomenon referred to as semantic priming (see Plaut, 2002, for an alternative feature-based model of semantic priming). Semantic priming successfully accounts for performance in a variety of tasks, such as lexical decision, sentence verification and word pronunciation (see McNamara, 2005; Neely, 2012 for reviews).

There is some evidence that semantic priming can extend to two or three steps within a network (e.g., lion – tiger – stripes). For example, Balota and Lorch (1986) used a mediated priming pronunciation task to show that response latencies to pronounce a target word (e.g., stripes) were faster following a directly related prime (e.g., tiger), which were faster than a mediated prime (e.g., lion), which in turn were faster than an unrelated prime (e.g., sand). McNamara and Altarriba (1988) extended this work and provided evidence for multiple-step priming in a lexical decision task. Importantly, in each of these studies, semantic steps within a network were not based on any a-priori model of semantic memory but based on items selected by the experimenters that appeared to have no *direct* relationship for the mediated pairs (e.g., lion and stripes), but did have direct relationships with the mediators (e.g., lion to tiger and tiger to stripes). More recently, Jones and Mewhort (2007) showed that mediated pairs such as those discussed above are in fact closer in a computational semantic space within a random vector accumulation model (i.e., BEAGLE). Thus, although mediated or two-step priming effects have

been obtained and afforded considerable theoretical discussion, there has been relatively little work investigating priming effects for more distant relationships, based on a priori defined network configurations, as opposed to items selected by the researchers' intuitions.

Understanding priming effects for more distant concepts requires an independent specification of the underlying network representation of these concepts. Recent graph theoretic approaches to knowledge representation afford a class of semantic network models, which represent words as nodes in a large memory network, where words with similar meanings are connected to each other via edges (see Steyvers & Tenenbaum, 2005). This approach of representing memory structures using graph theoretical methods is being increasingly used to study the large-scale structure of language and memory. For example, Steyvers and Tenenbaum (2005) constructed three types of semantic networks based on the Nelson association norms (Nelson, McEvoy & Schreiber, 2004), WordNet (Fellbaum, 1998; Miller, 1995) and Roget's thesaurus (Roget, 1911), and showed that each semantic network followed a small-world structure (Barabási & Albert, 1999; Strogatz, 2001), similar to several naturally occurring complex networks such as the World Wide Web (Albert, Jeong, & Barabási, 2000; Barabási & Albert, 1999; Watts & Strogatz, 1998). Additionally, Steyvers and Tenenbaum (2005) proposed a mechanism for language acquisition based on semantic growth and preferential attachment, which proposed that new concepts attach to already existing concepts that have more connections with other concepts in the network. Graph theoretical methods-based research has also been extended to speech production and lexical retrieval (Chan & Vitevitch, 2010; Vitevitch, Chan & Goldstein, 2014), creativity (Kenett, Anaki & Faust, 2014), memory retrieval (Vitevitch, Chan & Roodenrys, 2012), and similarity judgments (De Deyne, Perfors, & Navarro, 2016).

De Deyne, Navarro, and Perfors (2016) have recently examined how individuals assess similarities between weakly related words using a semantic network approach. In their study, participants were presented with word triads (e.g., *butter*, *train* and *saddle*) and asked to indicate which two words were most related. Their results indicated that there was systematic consistency in participant responses to seemingly unrelated word triads (e.g., most participants indicated that *train* and *saddle* were more related). Further, De Deyne et al. provided evidence that these similarities were most successfully captured through a spreading activation mechanism operating over a word association network, compared to other similarity measures based on local neighbors. Importantly, this work introduced a novel method of assessing similarity between distant concepts. However, to our knowledge, there is relatively little work examining the extent to which such network-based representations account for semantic priming performance, the most widely studied paradigm to examine semantic representation and processes.

In a particularly relevant and important study, Kenett, Levi, Anaki, and Faust (2017) recently used a semantic relatedness task to explore the impact of semantic network path length derived from an 800-word Hebrew mental lexicon on priming. The Hebrew Association-Correlation Network (ACN) was created using graph theoretical methods and correlations derived from continuous free-association responses of 60 participants to 800 target words (for complete methodology, see Kenett, Kenett, Ben-Jacob & Faust, 2011). This type of representation (described further below) combines both distant network connections as well as semantic correlations, and thus represents a more hybrid model of semantic memory. In order to test the viability of this network structure, Kenett et al. (2017) had participants make relatedness judgments for word pairs chosen from this Hebrew network with varying path lengths. They found that as network path length between the word pairs increased, fewer word pairs were

judged as related. Importantly, they also reported a quadratic relationship between network path length and response times (RTs) to make relatedness judgments, such that RTs increased for word pairs at shorter path lengths, but after path length 3, RTs systematically decreased for longer path lengths. This quadratic pattern likely reflected demands of the relatedness judgment task. Specifically, path length 1 was likely more clearly related than path length 2, while path length 3 was most ambiguous and hence produced the slowest response latencies. Importantly, Kenett et al. reported significant differences in RTs for items that were 4- and 6-steps apart. This pattern suggests that priming can potentially extend to relatively distant connections, i.e., to path lengths 4 and beyond for word pairs that are consistently judged as unrelated. They also showed that this network outperformed a popular distributed space model, Latent Semantic Analysis (LSA; Deerwester, Dumais, Furnas, Landauer, & Harshman, 1990; Landauer & Dumais, 1997) in explaining task performance. However, given that Kenett et al. used an association-correlation methodology based on a relatively small network of Hebrew words, it remains unknown whether other association networks (e.g., Steyvers & Tenenbaum, 2005) could also capture such distant semantic relationships at similar levels as the ACN. Moreover, it is important to extend the Kenett et al. network structure to a larger English-based network analysis to examine the generalizability of their findings.

The present set of experiments were designed to address three specific questions. First, we were interested in examining the extent to which the patterns of multiple-step priming reported by Kenett et al. in the Hebrew semantic distance task would replicate in a much larger semantic network in English, using the Kenett et al. network structure. We created this network from a large 5018-word database of free association norms collected by Nelson et al. (2004) and examined the extent to which path lengths predict task performance, after controlling for lexical

variables such as word frequency, length, concreteness and lexical decision times extracted from the English Lexicon Project (ELP; Balota et al., 2007). In our first experiment, following Kenett et al.'s procedure, participants were briefly presented a prime word for 120 ms, and then a target word for a relatedness decision. We examined whether network path length between the prime and target word predicted the extent to which a word pair was judged as related or unrelated, and whether the response latencies to make these judgments varied as a function of path length. If semantic network parameters from the ACN indeed predict performance, we should see an influence of path length on response latencies and replicate the quadratic relationship between path length and RTs, as described by Kenett et al. (2017).

Second, it is important to extend the Kenett et al. results to a different experimental paradigm that does not specifically direct attention to semantics, as in relatedness judgments. Specifically, a potential concern regarding semantic networks created through human association norms is a type of circularity, i.e., the success of association networks in explaining relatedness data could be due to shared variance with the task. In particular, Jones, Hills, and Todd (2012) have argued that responses in a free association task are an outcome of a retrieval operation on an underlying semantic representation, and their success in behavioral tasks may simply reflect the similarity between the experimental task (e.g., verbal fluency or relatedness judgements) and the free association task itself. Thus, it is possible that the quadratic relationship observed in Kenett et al. (2017) may reflect the specific nature of the relatedness judgment task, i.e., as noted above, the quadratic relationship likely reflects differences in how the “distance” between any two words might influence the “related” vs “unrelated” decisions in this task, and how individuals partition items into the two categories. We addressed this concern by employing a task that does not demand access to semantic information in an explicit manner. Thus, in

Experiment 2, participants first viewed a briefly presented prime (120 ms) and then identified targets through progressive demasking, a task that does not require directly attending to the relationship between the prime and target, and therefore removes the arbitrary distinction between related and unrelated word pairs. If network path length is indeed a measure of multiple-step semantic priming, we would expect to see response latencies to identify the target increase as the path length between prime and target words increases.

Finally, as noted above, we explored whether Step Distance Networks (SDNs) can also successfully capture semantic priming between relatively distant concepts. Given that the Kenett et al. network methodology involves computing correlations between word associations, it possibly captures more indirect associations between the words, unlike the direct associations captured by SDNs, such as those described by Steyvers and Tenenbaum (2005). In this light, one might consider the Kenett et al. structure to be a hybrid of network structure and latent structure. Thus, we compared the relative performance of two SDNs (Undirected and Directed) and the hybrid ACN created using the Kenett et al. methodology in Experiments 1 and 2. In the Directed SDN, words *a* and *b* were connected with an edge only if the word *a* evoked the word *b* as a response during free association, whereas in the Undirected SDN, words *a* and *b* were connected independent of the associative direction. In this way, we were able to examine the relative predictive power of each of these different network configurations in capturing any observed priming effects in two different behavioral tasks.

Experiment 1

Method

Network Construction. To construct the semantic networks, we used the 5018-word database of free-association norms collected by Nelson et al. (2004), in which 150 participants

on average wrote down the first word that came to mind in response to approximately 120 word-cues. We constructed three networks from this database: the ACN, Undirected SDN and the Directed SDN.

Association-Correlation Network (ACN). The ACN was created based on the methodology described by Kenett et al. (2011). Associative responses to 5018 cue words were first converted into a matrix, in which each column represented a cue word, and each row indicated unique associative responses for the target words. This matrix was converted to an association-correlation matrix, where the correlations between two target word profiles (i.e., the words produced to the two targets) was calculated based on the Pearson's formula. This correlation matrix was converted into a weighted, undirected network, such that each target word was a node in the network, and the correlation between two target words represented the weight of the edge between them. This fully connected network was then reduced to a Planar Maximally Filtered Graph (PMFG; Tumminello, Aste, Matteo, & Mantegna, 2005). The PMFG algorithm is an information filtering approach used to control for spurious correlations in correlation-based networks. PMFG draws edges between nodes by first sorting all the correlations in descending order, and only adding those edges to the graph that allow the resulting network to be embedded onto a sphere. This forms a planar network (a network in which no edges cross each other) with the same nodes but only those edges that represent the most relevant associations between target words. This unbiased topological constraint preserves more information compared to other network filtering approaches like the Minimum Spanning Tree (Tumminello et al., 2005) and has been applied to study semantic memory structure in clinical (Christensen, Kenett, Aste, Silvia, & Kwapil, 2018), and non-clinical populations (Borodkin, Kenett, Faust, & Mashal, 2016; Kenett et al., 2014). Path lengths between word-pairs were then calculated as the shortest path from one

word to another in this smaller PMFG network. Figure 1 (Top panel) displays a large-scale visualization of the ACN, and Figure 2 (Left panel) displays the 6-step shortest path from RELEASE to ANCHOR.

Undirected and Directed Step Distance Networks (SDNs). Following Steyvers and Tenenbaum (2005), in the Directed SDN, two words (a and b) were connected by an edge if the word a evoked the word b as an associative response for at least two participants in the Nelson database. If there was no directed association between two words in the Directed SDN, there was no edge. In the Undirected SDN, words were connected if the words were produced in response to each other, independent of the associative direction. Path lengths for each word pair in the network was calculated as the shortest path from one word to another. Figures 1 (Middle and bottom panels) and 2 (Middle and right panels) display visualizations of the two SDNs, and the shortest paths from RELEASE to ANCHOR.

Network Comparisons. Table 1 displays the network parameters for the three networks. As is evident from the large-scale visualizations, the ACN is sparser than the SDNs, with a greater clustering coefficient (an index of network connectivity, i.e., the extent to which neighborhoods of neighboring nodes overlap) and longer average path lengths, indicating more indirect, conceptual associations compared to the direct associations captured by SDNs with shorter path lengths overall.

Table 2 displays the correlation among the path lengths derived from each of the networks for the sets of words used in our experiments. As is clear from these correlations, there are considerable differences across the different types of network configurations (as displayed in Figure 1) and the relationships they capture (as shown in Figure 2). As shown in Figure 1, the correlation-based method used to create the ACN leads to a very sparsely connected network, in

which obscure, higher-level associations are closely represented (e.g., TRAGEDY-REMORSE is 1 step away), whereas several direct (e.g., VOLCANO-ASH is 15 steps away) and mediated associations (e.g., LION-STRIPES is 38 steps away) appear to have exaggerated distances. It is likely that the planarity criterion imposed during network construction for the ACN causes several direct associations to be dropped, primarily retaining indirect, higher-level conceptual relationships, an issue we return to in the General Discussion. This also produces the irregular shape of the network, as shown in Figure 1, where higher-level conceptual representations produce the branch-like structure, a pattern also displayed in the Kenett et al. model of Hebrew word associations. Additionally, we also see some path-based differences between the Undirected and Directed SDNs. Specifically, the Directed SDN had slightly longer paths compared to the Undirected SDN, which suggests that the direction of association may be important in tasks that involve forward or backward association. Overall, however, path lengths derived from the two SDNs were very highly correlated, suggesting that the directed and undirected associative networks largely overlap in their network structure, and differ from the ACN.

Participants. Forty Amazon Mechanical Turk users ($M_{age} = 36$ years, $SD = 11.3$) were recruited online and an additional 40 undergraduate students ($M_{age} = 20$ years, $SD = 1.2$) were recruited from Washington University in St. Louis. Amazon Mechanical Turk users were paid \$3.75 for their participation, and Washington University students received course credit for participation. Mean score on the Shipley Vocabulary Test was 31.62 ($SD = 6.69$) for the Mechanical Turk Users and 30.78 ($SD = 3.65$) for the Washington University students. Mean years of education was 14.68 ($SD = 2.72$) for the Mechanical Turk Users and 14 ($SD = 1.36$) for the Washington University students. All except 2 participants were self-reported native English

speakers, and their performance on the task did not differ from the group average and thus the final sample included all 80 participants. This and the following experiment were approved by the Institutional Review Board at Washington University in St Louis.

Materials. Because we initially wanted to extend and replicate the Kenett et al. study, we followed their procedure and randomly sampled forty word-pairs from path lengths 1, 2, 3, 4, 6 and 15 from the ACN. Although Kenett et al. used a single list, in order to increase generalizability, we created five different lists, with each list created using the same procedure. The stimuli thus consisted of 1200 distinct word-pairs across the 5 lists. The primes and targets in these lists were then counterbalanced across participants for order of presentation. For each word-pair sampled from the ACN, we also obtained corresponding path lengths in the Undirected and Directed SDNs.

Procedure

The relatedness task was developed using JSPsych, an online software for conducting psychological experiments. Each participant received a link to the experiment and completed the experiment online. Participants received task instructions and were guided through ten practice trials. We also included 15 buffer trials before the actual experiment, which were removed from final analyses. Following Kenett et al., on each trial, participants saw a fixation cross for 200 milliseconds, followed by a blank screen for 100 ms. The prime was then briefly presented for 120 ms, followed by the target for 120 ms. Participants decided whether the prime and target were related or unrelated and indicated their response by pressing a button (key “K” or “L” on the keyboard). The order of using “K” for related and “L” for unrelated was counterbalanced across participants. After responding, participants saw a blank screen for 500 ms before proceeding to the next trial.

Results

There were no significant differences in overall patterns for the sample recruited from Amazon Mechanical Turk, and the sample recruited from Washington University in St Louis. Further, the specific lists also did not influence the overall patterns. Therefore, all reported results contain the full sample of 80 participants across all lists.

Effect of ACN Path Length on RTs. To minimize the undue influence of extremely fast or slow RTs in our analyses, each individual's RTs were screened in the following manner for all analyses. First, RTs faster than 250 ms and slower than 2000 ms were removed. Second, a mean and standard deviation were calculated from the remaining trials for each participant and any RTs that exceeded 3 standard deviations (SDs) from the participant mean were also removed. This process excluded 5.4% of the total trials. After this trimming procedure, we standardized the remaining trials within each participant and conducted all primary analyses using trial-level standardized RTs, to minimize any effects of general slowing and individual differences across participants (see Faust, Balota, Spieler & Ferraro, 1999).

In order to exactly replicate the analyses procedures reported by Kenett et al., each of the path lengths were either classified as related or unrelated based on the percentage of related and unrelated responses to the specific word-pairs in each path length. Figure 4 (Top panel) displays percentage of related and unrelated responses in each path length for the ACN as well as the Kenett et al. data (Experiment 2). As shown, the present results closely replicated the patterns reported in Kenett et al. (2017), although it is important to note that the difference in proportions between the related and unrelated pairs at different path lengths were less extreme in our study, compared to Kenett et al., an issue we return to in the Discussion. Based on these percentages, and the criteria of at least 50% of words producing a related response, only word-pairs

corresponding to path length 1 were classified as related, and all other path lengths were classified as unrelated.

A path length (1, 2, 3, 4, 6, 15) repeated measures Analysis of Variance (ANOVA) was conducted to examine the effect of path length on mean standardized RT at the participant level. Following Kenett et al., only successful trials were analyzed, based on the a priori and post-priori classification of related or unrelated paths, described above. We observed a significant main effect of path length, $F(5, 395) = 7.42, p < .001, \eta_p^2 = .09$. Post hoc analyses revealed that RTs significantly increased from path length 1 to 2 ($p = .006$), and then decreased from path lengths 2 to 6 ($p = .001$). RTs also significantly decreased from path lengths 4 to 15 ($p = .015$). No significant differences between RTs were found when path length increased from 3 to 4, or 4 to 6, or 6 to 15. As shown in Figure 4 (Bottom panel), we successfully replicated the pattern reported by Kenett et al. for RTs as a function of path length in the ACN.

Next, to avoid the exclusion of trials based on whether they were judged as related or unrelated at different path lengths, we included response type (whether an item was judged as related or unrelated) as a predictor in a linear mixed effects (LME) analysis. We also controlled for lexical variables such as word frequency, length, concreteness and standardized lexical decision times by performing the LME with random intercepts at the participant, item and trial level. Lexical characteristics for each of the words in the experiment were obtained from the ELP (Balota et al., 2007), and mean frequency, length, concreteness and standardized lexical decision times for each word-pair were computed and included as covariates in our analyses. Importantly, we observed a significant interaction between path length and response type (related or unrelated), after controlling for item characteristics ($\Delta AIC = 34$ for model with and without interaction term, $p < .001$). Specific contrasts also revealed a significant difference in

standardized RTs between path lengths 1-related and 2-unrelated ($p < .001$), and also a significant difference between 6-unrelated and 15-unrelated ($p = .013$). Figure 5 displays the standardized RTs as a function of network path lengths for successful trials. Additionally, we also specifically tested for the presence of a quadratic trend in the RTs and found that the quadratic model was significant and explained more variance than the linear model ($\Delta AIC = 4, p = .014$). Thus, the quadratic trend persisted after controlling for type of relatedness decision as well as item-level differences. In order to ensure that using lexical decision latencies as a covariate was not potentially influencing the results, we also conducted an analysis without this covariate and the same pattern was observed. Figure 6 displays the full data for the relatedness decisions as a function of network path length in the ACN, because plotting only “successful” trials based on an arbitrary threshold may mask important differences in the pattern of RTs for “related” and “unrelated” items at the different path lengths. Interestingly, RTs for “related” word pairs systematically increased with longer path lengths, indicating slower processing for these distant items.

Effect of Step Distance Network Path Lengths on RTs. In addition to the ACN, as noted, we also created two SDNs based on the method used in Steyvers and Tenenbaum (2005). We examined the effect of path lengths derived from the Undirected and Directed SDNs on standardized RTs in the relatedness task. Because the SDNs had a more compact and densely connected representation (as seen in Figure 1), path lengths ranged from only 1-5 in the Undirected SDN, and from 1-8 in the Directed SDN. Due to extremely few data points in the higher path lengths in the directed network, we collapsed all items at path lengths greater than 5 (i.e., 120 in PL6, 19 in PL7, 4 in PL8, and 97 items with no paths) into one path length in our analyses. As shown in Figure 5, both the Undirected and Directed SDNs also showed a quadratic

trend for standardized RTs as a function of path length, with RTs significantly rising from path lengths 1 to 2 (p 's $< .001$) and then reliably decreasing from path length 2 onwards. Importantly, we observed a significant decline in RTs from path lengths 3-unrelated to 4-unrelated, in the Undirected ($p < .001$) and, and from path lengths 2-unrelated to 5-unrelated in the Directed SDN ($p < .001$). Figure 6 also shows the full data for the relatedness decisions as a function of SDN path lengths.

Discussion

The results from Experiment 1 provide strong evidence for distant priming in the relatedness judgment task, and also demonstrate a quadratic relationship between network path length and RTs to judge a word-pair from the network as related or unrelated. As network path length between the words increased, more word-pairs were judged as unrelated, and response latencies first increased until path length 2, and then systematically decreased for more distant path lengths. These results replicate and extend the pattern observed by Kenett et al. (2017) for the Hebrew network. Importantly, we found significant differences between RTs at distant path lengths, specifically path lengths 6 and 15 in the ACN, after controlling for item-specific lexical characteristics. This suggests that one can observe priming across quite distant relationships, within this paradigm.

There were some notable differences between our findings and those of Kenett et al. (2017). First, as shown in Figure 4 (Top panel), the difference in the proportion of participants who judged the word-pairs as related or unrelated at different path lengths was less extreme for our English word-pairs, compared to the Hebrew word-pairs used by Kenett et al. (2017). It is important to note that the word-pairs in our study were randomly sampled across 5 lists from the ACN, whereas Kenett et al. sampled only one set of 240 items for their study and matched the

items for length, frequency and concreteness. In contrast, we used these variables as covariates in our analyses, to reduce the possibility of potential item-selection effects (see Forster, 2000). Second, we found that word-pairs at path length 2 were judged as related and unrelated by a statistically equivalent proportion of participants, and the response latencies to make relatedness judgments were slowest for path length 2, compared to other path lengths in the network. Of course, this suggests that these items at path length 2 were most ambiguous regarding relatedness and hence produced the slowest response latencies. Kenett et al. observed the slowest RTs for path length 3 and argued that the breadth of spreading activation is at least 3 steps. Our findings differ in this regard, as we not only find that significant differences at shorter path lengths, i.e., between 1 and 2, but more importantly capture differences between even longer path lengths (i.e., 6 and 15) in the network, compared to Kenett et al. where they found significant differences only between path lengths 4 and 6, and not beyond.

In addition to replicating the pattern observed by Kenett et al. (2017), we provide strong evidence that directional and nondirectional SDNs can also capture similar distant semantic relationships between concepts. However, it is important to acknowledge that the distribution of word-pairs corresponding to each of the undirected and directed path lengths was not the same, because our sample was created to ensure equal number of items in each path length specifically for the ACN. Further, as shown in Figure 7, “distant” items in the ACN (i.e., path lengths 6 and 15) did not consistently correspond to distant items in the SDNs, i.e., over 50% of the items in each ACN path length did not correspond to the same path lengths in the Undirected and Directed SDNs. This is noteworthy, as it indicates that the ACN potentially exaggerates distances between word pairs, that may not be as many “steps” apart, at least based on step distance associative networks. Of course, this is not to say that these distant relationships are

unimportant. Having said this, we did find that path lengths in the Undirected and Directed SDNs predicted RTs, suggesting that step distance networks are also able to effectively account for the distant priming effects in the memory network.

It is important to reiterate that the relatedness decisions may be driving some of the observed priming effects in this task. First, the quadratic trends are likely due to the influence of the strength of relatedness on related and unrelated decisions, i.e., the reason one finds the slowest RTs at path length 2 is because these items are most ambiguous regarding their status as related word pairs. Importantly, due to the arbitrary nature of the relatedness decision, and the differences in proportion of related and unrelated items at different path lengths within the three networks, it is difficult to directly compare the relative performance of the three networks in this task. Second, the SDNs used in the current study were explicitly created from free association norms, and thus their explanatory power may just reflect the high degree of overlap between the base task (free association) and the relatedness judgment task used in Experiment 1. Thus, in Experiment 2, we explored whether network path length can indeed account for semantic priming in a priming task that does not explicitly demand direct access to the association to make the response (i.e., target demasking), and also compared the relative performance of each of the three network models.

Experiment 2

Methods

Participants. Forty young adults ($M_{age} = 20.9$ years, $SD = 2.8$) were recruited from undergraduate courses at Washington University in St Louis, and from Volunteers for Health (VFH), a recruitment program sponsored by the Washington University School of Medicine. We decreased the total number of participants here because now all trials involve the same response

within a participant and contribute to the analyses. All participants were Native English speakers and were compensated through course credit or \$10 for their participation. One participant misunderstood the experiment instructions and typed primes instead of targets for all trials, and hence their data was excluded from the final sample.

Materials. Given that the specific list did not influence any of the results in Experiment 1, one list of 240 items was randomly chosen from one of the five lists used in Experiment 1, with 40 word-pairs from path lengths 1, 2, 3, 4, 6 and 15 randomly sampled from the ACN. The list was then counterbalanced for directionality between primes and targets across subjects. Each word pair also had a corresponding path length in the Undirected and Directed SDN.

Procedure. The primed progressive demasking task was developed using E-Prime 2.2. Participants saw a black fixation cross on the screen for 500 ms. Next, a blank screen was displayed for 200 ms, followed by the prime word (e.g., RELEASE), which was displayed for 120 ms. Immediately after, the target word was progressively demasked on the screen at the same location (see Figure 3). During progressive demasking, the display alternated between the target (e.g., ANCHOR) and a mask (a row of pound signs matching the length of the word, e.g., #####). The total duration of target-mask pair was held constant at 500 ms but the ratio of target display time to mask display time progressively increased. In the first cycle, the mask was presented for 500 ms. In the second cycle, the target was displayed for 16 ms followed by the mask for 484 ms. The duration of the target increased at each cycle (0, 16, 32,...,500 ms) and the duration of the mask decreased (500, 484, 468,...,0 ms). The demasking procedure continued until the target was fully revealed for 500 ms, or until the target was identified by the participants by pressing the spacebar. Participants then typed in the correct target word on the next screen. The next trial began immediately after typing in the correct target and pressing spacebar.

Participants were given 3 practice trials, followed by 240 experimental trials. After every 18 trials, participants could take a short break and continued with the experiment when they were ready.

Results

Effect of ACN Path Length on RTs. Before analyzing the response latencies to identify the target words, we first removed all trials in which the participant did not identify the correct target, which excluded 2.7% of the total trials. Next, we excluded outliers and standardized the RTs using the same procedures as in Experiment 1. This process excluded 1.9% of the remaining trials. A repeated measures ANOVA on standardized RTs revealed a significant effect of ACN path length, $F(5,190) = 53.85$, $p < 0.001$, $\eta_p^2 = .586$. As shown in Figure 8, this effect indicated a significant increase in RTs from path lengths 1 to 2 ($p < .001$), and 2 to 3 ($p < .001$). Differences between RTs to identify the target at path length 3 and higher path lengths were not statistically significant. Importantly, these effects persisted after controlling for lexical variables such as word frequency, length, concreteness and standardized lexical decision times in LME analyses, as in Experiment 1.

Effect of SDN Path Lengths on RTs. We examined the effect of path lengths derived from the Undirected and Directed SDNs on standardized RTs in the primed progressive demasking task. Due to extremely few data points in the higher path lengths in the directed network, we again collapsed all path lengths greater than 5 into one path length, as in Experiment 1. As shown in Figure 8, path lengths from the Undirected SDN significantly predicted RTs to identify the target. Specific comparisons indicated that RTs increased from path length 1 to 2 ($p = .001$), from path lengths 2 to 3 ($p < .001$), and then marginally from path lengths 3 to 4 ($p = .058$) in the Undirected SDN. Path lengths from the Directed SDN also significantly predicted

RTs to identify the target. Specific comparisons indicated that response times significantly increased from path lengths 2 to 3 ($p = .015$), from 4 to 5 ($p = .038$), and then from path length 5 to higher path lengths ($p = .001$) in the Directed SDN.

Model Comparisons. Given that the results from this task were not complicated by the relatedness decision (i.e., response latencies should be linearly related to demasking performance), we were able to directly compare the model estimates. To estimate the unique variance contributed by the distance estimates derived from each type of network configuration at the item level, we calculated the individual R^2 for each model, as well as estimates of AIC (Akaike, 1987) and BIC (Schwarz, 1978), after controlling for covariates. Importantly, we used continuous distance estimates for these comparisons to retain the numerical distance-based differences among the different network models and examine how the specific “steps” in the networks explain demasking latencies at the item level. Following Wagenmakers and Farrell (2004), we also calculated Akaike and Schwarz weights for each network to assess the relative strength of evidence for each model based on their AIC and BIC model estimates. As shown in Table 3, the Directed SDN was the most likely model, followed by the Undirected SDN, which in turn was better than the ACN, based on Akaike and Schwarz weights. Further, all three network models explained a significant amount of variance over and above the item-level covariates, i.e., word length, concreteness, word frequency, and lexical decision times.

Discussion

Results from Experiment 2 indicate that network path lengths can indeed account for distant semantic priming in a primed progressive demasking task. We found a linear relationship between network path length and response latencies to identify the target via demasking in the ACN, the Undirected and Directed SDNs. This is especially interesting as the progressive

demasking task did not require any direct retrieval of the semantic association to make the response, and yet, we observed that path lengths derived from word associations directly predicted demasking response latencies. Interestingly, path lengths from the ACN increased linearly only up to 3 steps, after which the network seemed to no longer be sensitive to priming effects in this task, further suggesting some differences in the sensitivity across the tasks. As expected, the rise in RTs in this task mirrored the rise in RTs for the “related” item decisions in Experiment 1 (see Figure 6), indicating that priming extends to distant concepts but also dissipates as the distance between the concepts increases.

General Discussion

The current set of experiments investigated the influence of large-scale word association networks on priming effects in two behavioral tasks. We provide strong evidence for multiple-step priming using network path length from association networks as an indicator of distance between concepts within the network. We now discuss specific findings from the experiments, and important differences between the different network representations.

Network Structure and Distant Priming

A primary goal of the present study was to empirically examine priming effects for distant connections, as defined by distinct network representations. This work was motivated by a recent paper by Kenett et al. (2017) who used network path length derived from an 800-word Hebrew network to show that path length predicted performance in a relatedness judgment and free recall task. Specifically, Kenett et al. found a quadratic relationship between network path length and response latencies in a relatedness judgment task, such that response latencies overall increased until a categorical boundary (i.e., 3 steps) in the network and then decreased at longer path lengths. We successfully replicated and extended their work to a larger 5018-word

association network in English and also compared their graph-theoretical approach of mapping the lexicon to undirected and directed step distance networks (Steyvers & Tenenbaum, 2005). Further, while Kenett et al. did not observe any differences for response latencies at distant path lengths (i.e., beyond 6 steps) we found significant decreases in response latencies at path lengths 6 and 15 in the ACN, and at path lengths 3, 4 and 5 in the SDNs, after controlling for lexical variables such as word frequency, concreteness, word length and lexical decision times. Our results thus provide clear evidence for distant semantic priming and add to previous work on multiple-step spreading activation (Balota & Lorch, 1986; Kenett et al., 2017; McNamara & Altarriba, 1988). To our knowledge, this is the first study to empirically demonstrate that semantic priming can indeed extend to relatively distant concepts in the network, i.e., 6 or 15 steps, within the ACN configuration. The results from the second experiment further indicated that network path length also successfully accounts for semantic priming in a task that does not demand direct retrieval of the association (as via a relatedness judgment task used in Kenett et al., and in our first experiment). We again found significant differences in response latencies to identify the target word at relatively distant path lengths, i.e., 4 and 5 steps. Further, response latencies did not differ after 3 steps in the ACN but continued to increase linearly in the Undirected and Directed SDNs. Based on overall model fits, the Directed SDN was the most likely model, followed by the ACN and the Undirected SDN. Importantly, all models significantly explained more variance than item-level covariates.

Comparing Word Association Network Configurations

An important contribution of the current set of experiments is the comparisons across three different network configurations. However, one potential concern regarding these comparisons may be that the ACN and SDNs may not be truly comparable, due to the

differences in their network construction methodologies. Specifically, the ACN starts from a correlation matrix for all words and applies the PMFG algorithm to construct the final network. On the other hand, the SDNs only connect edges between the nodes if at least two participants produced one word in response to another, and this criterion is directional for the Directed SDN. To evaluate whether these differences in methodologies influenced our findings, we constructed three new networks, i.e., an Undirected PMFG network, an Unfiltered ACN, and a Filtered ACN.1, and examined their predictive power in our experiments.

To construct the Undirected PMFG, we first calculated correlations between all nodes in the Undirected SDN using the undirected path length distance matrix. Next, the PMFG algorithm was applied to this correlation matrix to construct a planar graph consisting of the same number of nodes and retaining only the most relevant edges. As shown in Figure 9 (Top Panel), the Undirected PMFG had a very similar configuration to the ACN, with path lengths ranging from 1 to 25. To ensure an equal number of items in each “path length” for our experiments, we partitioned these original Undirected PMFG path lengths into quintiles and used them in our subsequent analyses. To construct the Unfiltered ACN, we simply retained the matrix of symmetric correlations that was later passed to the PMFG algorithm to create the original ACN. Thus, the Unfiltered ACN is a complete network with edges weighted by the correlations (see middle panel of Figure 9). Given that the Unfiltered ACN contains correlations between words and not “path lengths”, we partitioned these correlations into quintiles (arbitrary steps in this case) for the items used in our experiments. Finally, to construct the Filtered ACN.1 network, we dropped all edges with correlations below 0.1 in the Unfiltered ACN and then constructed an undirected, unweighted network (see bottom panel of Figure 9). Path lengths in the Filtered

ACN.1 ranged from 1 to 3 and were directly used as independent variables in our analyses. Network parameters for these new networks are presented in Table 2.

To evaluate the extent to which these new networks compare to our original networks, we repeated our analyses for Experiments 1 and 2. In Experiment 1, we found that the Undirected PMFG, Unfiltered ACN and Filtered ACN.1 all accounted for the quadratic pattern in relatedness judgement response latencies. More importantly, we also specifically estimated the relative variance accounted for by each network in response latencies to identify the target through demasking in Experiment 2. We first compared the non-planar networks, i.e., the Directed and Undirected SDNs, Unfiltered ACN, and the Filtered ACN.1. As shown in Table 4, we again found that the Directed SDN was the most likely model, compared to the Filtered ACN.1 network, which was in turn more likely than the Undirected SDN, followed by the Unfiltered ACN, based on Akaike and Schwarz weights. Interestingly, the Unfiltered ACN and Undirected SDN explained relatively equivalent amount of variance, suggesting that the inclusion of directional information in the Directed SDN contributes significantly to its predictive power in this task. On the other hand, when comparing the two planar networks (ACN and Undirected PMFG), the ACN was a more likely model, compared to the Undirected PMFG. Collectively, these findings suggest that alternate ways of constructing the ACN from similarity correlations (e.g., using an unfiltered network or setting an arbitrary cutoff) lead to a slight increase in explanatory power compared to the original ACN, possibly due to the retention of greater number of edges compared to the original ACN. However, the Directed SDN still explains the maximum variance in this task, likely due to the directed associations it captures. Further, applying the PMFG algorithm to an already restricted network such as the Undirected SDN actually leads to a loss in explanatory power compared to the ACN.

As described earlier, the ACN uses co-occurrence information and an information filtering algorithm to construct the network. This type of network construction method leads to several direct associations (e.g., TIGER-STRIPES is 37 steps away in the ACN and directly connected, i.e., 1 step away in the SDNs) being dropped, giving rise to more indirect, high-level associations (e.g., TRAGEDY-REMORSE is 1 step away in the ACN and farthest, i.e., 4 steps away in the Directed SDN). While this topological constraint of graph planarity represents an unbiased method for eliminating spurious correlations from the complete network, it is possible that imposing this criterion results in a loss of direct links in the network, thus losing the shortest paths between words and potentially exaggerating distances between some items. Indeed, this is exemplified in the branching network structure shown in the top panel of Figure 1. Previous work has shown how metric axioms that must be respected by spatial representations (Tversky, 1977) are routinely violated by word association norms (Griffiths, Steyvers, & Tenenbaum, 2007). Thus, it is possible that the planarity criterion used by the ACN is a similar geometric constraint that in some instances does not necessarily capture *direct* word associations. Furthermore, our additional analyses of the Unfiltered ACN, Filtered ACN.1 and Undirected PMFG provide evidence that when a different geometric criterion is used for network construction, the ACNs and SDNs may be differentially sensitive to explaining behavioral performance in priming tasks.

Finally, it is worth noting that the ACNs and SDNs may perform differently in a conceptually driven task where different types of semantic relationships are accessed, which would suggest that different types of stimuli/tasks emphasize different properties of the semantic network space. Indeed, Gruenenfelder, Recchia, Rubin, and Jones (2015) recently argued for a hybrid representation of lexical/semantic memory and suggested that individuals switch between

a contextual representation and associative networks when generating free associations. Our results suggest that there may also be differences in how individuals use these different types of semantic representations in tasks that do not explicitly involve word association but do place constraints on the type of semantic relationships being accessed. Our future work aims to address the extent to which these models predict performance in conceptual semantic tasks.

Distributional Models of Word Representation

While the current set of experiments focused on different types of network representations, another important class of models of semantic memory represent words through vectors in a multidimensional space. In distributional representations of semantic knowledge such as the LSA (Landauer & Dumais, 1997), BEAGLE (Jones & Mewhort, 2007) and *word2vec* (Mikolov, Chen, Corrado & Dean, 2013), words are an aggregate of distributed dimensions that are typically derived from statistical co-occurrences in natural language. This type of representation of semantic memory is clearly different from a network-based perspective, which are typically based on word association norms. Consequently, there has been considerable interest in comparing different types of semantic word representations and the extent to which they explain complex behavior. For example, LSA has been shown to successfully simulate complex human behavior in tasks such as word categorization (Laham, 2001), semantic similarity (Landauer & Dumais, 1997) and discourse comprehension (Kintsch, 1998). However, LSA has also had some difficulty accounting for semantic priming effects (Hutchison, Balota, Cortese & Watson, 2008; Kenett et al, 2017), ignoring word transitions in language (Perfetti, 1998), and violating power laws of semantic connectivity observed in real networks (Steyvers and Tenenbaum, 2005). To our knowledge, there is relatively little work examining the extent to which network and distributional representations in the English language account for semantic

priming performance, especially for more distant concepts. Although, as discussed before, Kenett et al. showed that the ACN path lengths in the Hebrew network outperformed LSA in the relatedness judgment task.

Therefore, we conducted additional analyses to examine the extent to which corpora-based semantic models compare to association network-based semantic models described above in explaining priming effects. As discussed before, Hutchison et al. (2008) previously showed that LSA does not account for semantic priming effects, compared to simple associative strength estimates. However, the Hutchison et al. study was conducted on direct associates and not distant model-based associates, and hence it a priori unclear if the LSA representation may do better in accounting for the present results.

In addition to comparing network-based models to LSA representations, we were also interested in testing an alternative distributional model. A more recent predictive distributional model, *word2vec* (Mikolov et al., 2013) has received considerable attention in fields of computer science and natural language processing for explaining performance in a variety of behavioral tasks. *word2vec* uses neural networks and a large training corpus (e.g., from a Google News dataset) to compute continuous vector representations of words which encode semantic information. These vector representations can then be used to compute an index of semantic similarity between words via vector cosines and are also useful inputs for other natural language processing tasks such as sentiment analysis (dos Santos & Gatti, 2014), document classification (Lilleberg, Zhu & Zhang, 2015), and named entity recognition (Severyn & Moschitti, 2015). Interestingly, *word2vec* has been shown to successfully solve verbal analogy problems (e.g., king : queen :: man:?) using simple vector arithmetic, although other research suggests that *word2vec* successfully captures only certain types of semantic relationships and not others (Chen, Peterson

& Griffiths, 2017). More recently, Mandera, Keuleers, and Brysbaert (2017) compared the relative performance of distributional semantic models (e.g., LSA-type and *word2vec*) on a battery of semantic priming tasks (Hutchison et al., 2013) and concluded that predictive distributional models like *word2vec* provided a better fit to the data. They also argued that predictive models are psychologically more plausible and computationally more compact than typical count-based distributional models like LSA (but see Goldberg, Levy, & Dagan, 2014 for a counterargument). However, when comparing *word2vec* to association-based network models, De Deyne, Perfors, and Navarro (2016) found word association networks consistently outperformed *word2vec* and count-based distributional models, even though the network models were trained on relatively smaller corpora. Of course, these studies do not shed light on the extent to which distributional models like *word2vec* and LSA explain priming performance for distant concepts, compared to different types of association-based semantic networks, which was the goal of these set of analyses.

In order to directly compare the different models, for all word pairs used in the current experiments, we obtained LSA cosines from the LSA website (<http://lsa.colorado.edu/>) using the recommended topic space of 300 factors, which corresponds to a general reading level (up to 1st year of college). We also obtained *word2vec* cosines from a pre-trained model trained on 100 billion words from a Google News dataset (Mikolov et al., 2013). Table 2 reports correlations between the vector cosines derived from LSA and *word2vec* models and the different networks for the stimuli used in Experiments 1 and 2. It is important to note here that there were considerable differences across the models in the extent to which they captured “semantic similarity”, given that the average correlation among all the different word representations across both experiments was only 0.46.

To estimate the variance contributed by each type of distance estimate, we computed separate estimates for R_2 , AIC and BIC for each network and distributional model in Experiment 2. Our results indicated that cosines derived from LSA ($R_2 = 24.97\%$, AIC = 505.11, BIC = 533.34) and *word2vec* ($R_2 = 27.72\%$, AIC = 489.52, BIC = 517.75) also successfully explained performance in the priming task. While the Directed SDN outperformed LSA, *word2vec* was the most likely model in this task overall. Our results are thus consistent with Mandera et al. (2017) in that the *word2vec* model does indeed explain distant priming effects in our task and has lower AIC/BIC values compared to all network-based models. Of course, continuous cosines from LSA and *word2vec* have more variability compared to the “steps” in the network models, so these comparisons are limited in scope, and how step-based representations derived from *word2vec* and LSA would compare to association networks is an avenue for future research.

Another important aspect of these results is the overall low correlations observed between the different distance estimates across distributional and network models. These correlations suggest that there are structural differences between network-based and distributional representations. For example, the word RELEASE is only 2 steps away from the word ANCHOR in the Undirected SDN but is very weakly associated in the *word2vec* (cosine -0.004) and LSA (cosine .08) multidimensional spaces. The path from RELEASE to ANCHOR is mediated by the word HOLD in the undirected network, but it is possible that this particular usage of ANCHOR does not co-occur in the same contexts as RELEASE in natural language, which is the mechanism underlying cosines obtained from the distributional models. Importantly, the tasks in the current study focused on semantic priming effects measured via progressive demasking, and it is possible that distributional representations may be more or less predictive of performance in more conceptual tasks.

The current results also inform an ongoing debate in semantic memory representation between association-based network models and distributional models. There is now accumulating evidence that distributional models that derive their semantic representations from solely linguistic sources are less likely to capture surface-level, attributive, and perceptual features (Baroni & Lenci, 2008; Lucy & Gauthier, 2017) and also encounter difficulties in explaining word association data (Griffiths, Steyvers, & Tenenbaum, 2007). Further, the amount of data required (e.g., billion words) to adequately train distributional models to perform at the level of association-based models calls into question their psychological plausibility (Asr, Willits, & Jones, 2014; De Deyne, Perfors, & Navarro, 2016). On the other hand, while association-based networks appear to perform at similar levels, and often outperform text-based distributional models in a variety of semantic tasks, the validity of such representations as complete accounts of semantic memory has been questioned on grounds of being constructed from retrieval-based processes involved in word association tasks (for a detailed discussion, see Jones, Hills, & Todd, 2012; Siew et al., 2018). However, there is evidence to suggest that association-based network models do indeed capture complementary semantic information compared to text-based distributional models (Gruenenfelder, Recchia, Rubin, & Jones, 2016). Therefore, a complete account of semantic memory should be able to account for how such associations are formed and acquire the complex network structure that successfully explains behavioral performance in semantic tasks. Ultimately, recent approaches that attempt to integrate non-linguistic information sources with traditional distributional models to construct multimodal semantic representations (Bruni, Tran, & Baroni, 2014; Kiela & Bottou, 2014; Lazaridou, Pham, & Baroni, 2015) appear to represent a promising step towards reconciling these two families of semantic models.

Limitations

There are a few important limitations to this work. First, even though the current results closely replicate the results reported by Kenett et al. (2017) in a Hebrew network, there is an important difference between the Hebrew ACN used by Kenett et al., and the English ACN in the current study. The Hebrew network was based on responses from a continuous free association task (where participants produced as many responses as they can to the target word), whereas the Nelson et al. norms are based on a discrete free association task (where participants produced the first word that came to mind for a particular target word). There is some debate regarding the validity of both continuous responses (see Nelson, McEvoy & Dennis, 2000) and discrete responses (Hahn, 2008). However, given that the English ACN and SDN networks used in the current set of experiments were created from the same Nelson et al. norms from a discrete association task, we believe that the differences observed in predictive power of the ACN and the SDNs in the current study were not critically influenced by the nature of associative responses per se, although it is important to acknowledge that this issue is deserving of further exploration.

Further, there were also important differences between the distributional (i.e., LSA and *word2vec*) and network-based models (i.e., ACN, Directed and Undirected SDN) of word representation. For example, the *word2vec* model used in the current set of experiments was trained on a Google News corpus, whereas the LSA model was derived from a pre-existing topic space intended to simulate general college-level reading levels. These corpuses are clearly very different from each other, and also from the Nelson et al. database of free association norms, and previous research suggests that the type of corpus used can significantly impact how well semantic models account for human performance (Recchia & Jones, 2009). Thus, it is important

to acknowledge that the nature of the task, the stimuli and the training corpora are all likely to influence the extent to which different types of semantic models explain behavioral performance.

Conclusion

The current set of experiments investigated the predictive power of path lengths derived from three large-scale semantic networks in accounting for lexical priming effects in two behavioral tasks and provided strong evidence for distant priming effects. We also demonstrated important structural differences between correlation-based networks and association networks and showed that association networks proposed by Steyvers and Tenenbaum (2005) are also able to capture relatively distant semantic relationships. Finally, we showed that distributional models like the LSA and *word2vec* also successfully captured similar behavioral patterns across the two tasks, although there were again important structural differences in the semantic information captured across the different models.

References

- Akaike, H. (1987). Factor analysis and AIC. In *Selected Papers of Hirotugu Akaike* (pp. 371-386). Springer, New York, NY.
- Albert, R., Jeong, H., & Barabási, A. L. (2000). Error and attack tolerance of complex networks. *Nature*, *406*(6794), 378.
- Anderson, J. R. (2013). *The Architecture of Cognition*. Cambridge, MA: Harvard University Press.
- Balota, D. A., & Lorch, R. F. (1986). Depth of automatic spreading activation: Mediated priming effects in pronunciation but not in lexical decision. *Journal of Experimental Psychology: Learning, memory, and cognition*, *12*(3), 336.
- Balota, D. A., Yap, M. J., Hutchison, K. A., Cortese, M. J., Kessler, B., Loftis, B., ... & Treiman, R. (2007). The English lexicon project. *Behavior Research Methods*, *39*(3), 445-459.
- Barabási, A. L., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, *286*(5439), 509-512.
- Botvinick, M. M., & Plaut, D. C. (2006). Short-term memory for serial order: a recurrent neural network model. *Psychological Review*, *113*(2), 201.
- Chan, K. Y., & Vitevitch, M. S. (2010). Network structure influences speech production. *Cognitive Science*, *34*(4), 685-697.
- Chen, D., Peterson, J. C., & Griffiths, T. L. (2017). Evaluating vector-space models of analogy. *arXiv preprint arXiv:1705.04416*.
- Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, *82*(6), 407.

- Collins, A. M., & Quillian, M. R. (1969). Retrieval time from semantic memory. *Journal of Verbal Learning and Verbal Behavior*, 8(2), 240-247.
- Deerwester, S., Dumais, S. T., Furnas, G. W., Landauer, T. K., & Harshman, R. (1990). Indexing by latent semantic analysis. *Journal of the American Society for Information Science*, 41(6), 391-407.
- dos Santos, C., & Gatti, M. (2014). Deep convolutional neural networks for sentiment analysis of short texts. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers* (pp. 69-78).
- Faust, M. E., Balota, D. A., Spieler, D. H., & Ferraro, F. R. (1999). Individual differences in information-processing rate and amount: implications for group differences in response latency. *Psychological Bulletin*, 125(6), 777.
- Fellbaum, C. (1998). Towards a representation of idioms in WordNet. *Usage of WordNet in Natural Language Processing Systems*.
- Gruenenfelder, T. M., Recchia, G., Rubin, T., & Jones, M. N. (2016). Graph-theoretic properties of networks based on word association norms: implications for models of lexical semantic memory. *Cognitive Science*, 40(6), 1460-1495.
- Hahn, L. W. (2008). Overcoming the limitations of single-response free associations. *Electronic Journal of Integrative Biosciences*, 5(1), 25-36.
- Hutchison, K. A., Balota, D. A., Cortese, M. J., & Watson, J. M. (2008). Predicting semantic priming at the item level. *The Quarterly Journal of Experimental Psychology*, 61(7), 1036-1066.

- Hutchison, K. A., Balota, D. A., Neely, J. H., Cortese, M. J., Cohen-Shikora, E. R., Tse, C.-S., ... Buchanan, E. (2013). The semantic priming project. *Behavior Research Methods*, 45(4), 1099–1114. <http://doi.org/10.3758/s13428-012-0304-z>
- Jones, M., Gruenfelder, T., & Recchia, G. (2011). In defense of spatial models of lexical semantics. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 33, No. 33).
- Jones, M. N., & Mewhort, D. J. (2007). Representing word meaning and order information in a composite holographic lexicon. *Psychological Review*, 114(1), 1.
- Kenett, Y. N., Anaki, D., & Faust, M. (2014). Investigating the structure of semantic networks in low and high creative persons. *Frontiers in Human Neuroscience*, 8, 407.
- Kenett, Y. N., Kenett, D. Y., Ben-Jacob, E., & Faust, M. (2011). Global and local features of semantic networks: Evidence from the Hebrew mental lexicon. *PLoS ONE*, 6(8), e23912.
- Kenett, Y. N., Levi, E., Anaki, D., & Faust, M. (2017). The semantic distance task: Quantifying semantic distance with semantic network path length. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 43(9), 1470.
- Kintsch, W. (1988). The role of knowledge in discourse comprehension: A construction-integration model. *Psychological Review*, 95(2), 163.
- Laham, R. D. (2000). Automated content assessment of text using latent semantic analysis to simulate human cognition (Doctoral dissertation, University of Colorado).
- Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, 104(2), 211.

- Lilleberg, J., Zhu, Y., & Zhang, Y. (2015, July). Support vector machines and word2vec for text classification with semantic features. In *Cognitive Informatics & Cognitive Computing (ICCI* CC), 2015 IEEE 14th International Conference on* (pp. 136-140). IEEE.
- Loftus, E. F. (1973). Activation of semantic memory. *The American Journal of Psychology*, 331-337.
- McNamara, T. P., & Altarriba, J. (1988). Depth of spreading activation revisited: Semantic mediated priming occurs in lexical decisions. *Journal of Memory and Language*, 27(5), 545.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems* (pp. 3111-3119).
- Miller, G. A. (1995). WordNet: a lexical database for English. *Communications of the ACM*, 38(11), 39-41.
- Neely, J. H. (1991). Semantic priming effects in visual word recognition. A selective review of current findings and theories. In D. Besner & G. Humphreys (Eds.), *Basic processes in reading: Visual word recognition*. Hillsdale, NJ: Lawrence Erlbaum.
- Nelson, D. L., McEvoy, C. L., & Dennis, S. (2000). What is free association and what does it measure?. *Memory & Cognition*, 28(6), 887-899.
- 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(3), 402-407.
- Perfetti, C. A. (1998) The limits of co-occurrence: Tools and theories in language research, *Discourse Processes*, 25:2-3, 363-377, DOI: [10.1080/01638539809545033](https://doi.org/10.1080/01638539809545033)

- Plaut, D. C. (2002). Graded modality-specific specialisation in semantics: A computational account of optic aphasia. *Cognitive Neuropsychology*, *19*(7), 603-639.
- Roget, P. M. (1911). *Roget's Thesaurus of English Words and Phrases...* TY Crowell Company.
- Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, *6*(2), 461-464.
- Severyn, A., & Moschitti, A. (2015, August). Learning to rank short text pairs with convolutional deep neural networks. In *Proceedings of the 38th international ACM SIGIR Conference on Research and Development in Information Retrieval* (pp. 373-382). ACM.
- Smith, E. E., Shoben, E. J., & Rips, L. J. (1974). Structure and process in semantic memory: A featural model for semantic decisions. *Psychological Review*, *81*(3), 214.
- Steyvers, M., & Tenenbaum, J. B. (2005). The large-scale structure of semantic networks: Statistical analyses and a model of semantic growth. *Cognitive Science*, *29*(1), 41-78.
- Strogatz, S. H. (2001). Exploring complex networks. *Nature*, *410*(6825), 268.
- Vitevitch, M. S., Chan, K. Y., & Goldstein, R. (2014). Insights into failed lexical retrieval from network science. *Cognitive Psychology*, *68*, 1-32.
- Vitevitch, M. S., Chan, K. Y., & Roodenrys, S. (2012). Complex network structure influences processing in long-term and short-term memory. *Journal of Memory and Language*, *67*(1), 30-44.
- Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of 'small-world' networks. *Nature*, *393*(6684), 440.

Footnotes

¹ Following Kenett et al. (2017), Figure 5 displays standardized RTs for only “correct” responses, i.e., RTs for related responses for path length 1 in all networks, and RTs for unrelated responses for all other path lengths.

Table 1

Summary statistics for semantic networks

Variable	Step Distance Networks (SDNs)			Association-Correlation Network (ACNs)			
	Undirected SDN	Undirected PMFG	Directed SDN	ACN PMFG	Unfiltered ACN	Filtered ACN.1	Hebrew ACN
n	5018	5018	5018	5018	5018	5018	800
$\langle k \rangle$	22	5.99	12.7	5.85	5018	95.19	5.94
L	3.04	12.5	4.27	23	1	2.45	10
D	5	27	10	61	1	4	25
C	.186	0.72	.186	.69	1	0.26	.68
L_{random}	3.03	1.95	4.26	1.95	3.03	--	3.94
C_{random}	.004	.05	.004	.05	.004	--	.005

Note. n = the number of nodes; $\langle k \rangle$ = the average number of connections; L = the average shortest path length; D = the diameter of the network; C = clustering coefficient; L_{random} = the average shortest path length with random graph of same size and density; C_{random} = the clustering coefficient for a random graph of same size and density

Table 2

Correlation matrix of network path lengths and word2vec cosines for items in Experiments 1 & 2

Experiment 1

	ACN	Undirected SDN	Directed SDN	LSA cosines	<i>word2vec</i> Cosines
ACN	1	-	-	-	-
Undirected SDN	.48	1	-	-	-
Directed SDN	.33	.57	1	-	-
LSA Cosines	-.28	-.45	-.39	1	-
<i>word2vec</i> Cosines	-.37	-.55	-.44	.59	1

Experiment 2

	ACN	Undirected SDN	Directed SDN	LSA Cosines	<i>word2vec</i> Cosines
ACN	1	-	-	-	-
Undirected SDN	.49	1	-	-	-
Directed SDN	.37	.58	1	-	-
LSA Cosines	-.36	-.51	-.41	1	-
<i>word2vec</i> Cosines	-.46	-.55	-.46	.64	1

Table 3

Model comparison metrics for the three networks in Experiment 2. Higher Akaike and Schwarz weights indicate greater evidence for the likelihood of a model.

Model	R ₂ (%)	AIC	Akaike Weights	BIC	Schwarz Weights
Covariates + Directed SDN	25.69	501.06	0.99	529.29	0.99
Covariates + Undirected SDN	21.80	522.32	2.41 x 10 ⁻⁵	550.55	2.4 x 10 ⁻⁵
Covariates + ACN	17.95	542.41	1.05 x 10 ⁻⁹	570.64	1.05 x 10 ⁻⁹
Covariates	13.06	564.55	1.64 x 10 ⁻¹⁴	588.74	1.23 x 10 ⁻¹³

Table 4

Model comparison metrics for planar and non-planar ACNs and SDNs. All models include item-level covariates. Higher Akaike and Schwarz weights indicate greater evidence for the likelihood of a model.

Model Type	Model	R ₂ (%)	AIC	Akaike Weights	BIC	Schwarz Weights
Non-planar networks	Directed SDN	25.69	501.06	0.99	529.29	0.99
	Filtered ACN.1	22.84	516.74	.0004	544.97	.0004
	Undirected SDN	21.80	522.32	2.41 x 10 ⁻⁵	550.55	2.41 x 10 ⁻⁵
	Unfiltered ACN	21.77	522.49	2.21 x 10 ⁻⁵	550.72	2.21 x 10 ⁻⁵
Planar networks	ACN	17.95	542.41	1.05 x 10 ⁻⁹	570.64	1.05 x 10 ⁻⁹
	Undirected PMFG	14.36	560.22	1.42 x 10 ⁻¹³	588.45	1.42 x 10 ⁻¹³
Covariate Model	Covariates	13.06	564.55	1.64 x 10 ⁻¹⁴	588.74	1.22 x 10 ⁻¹³

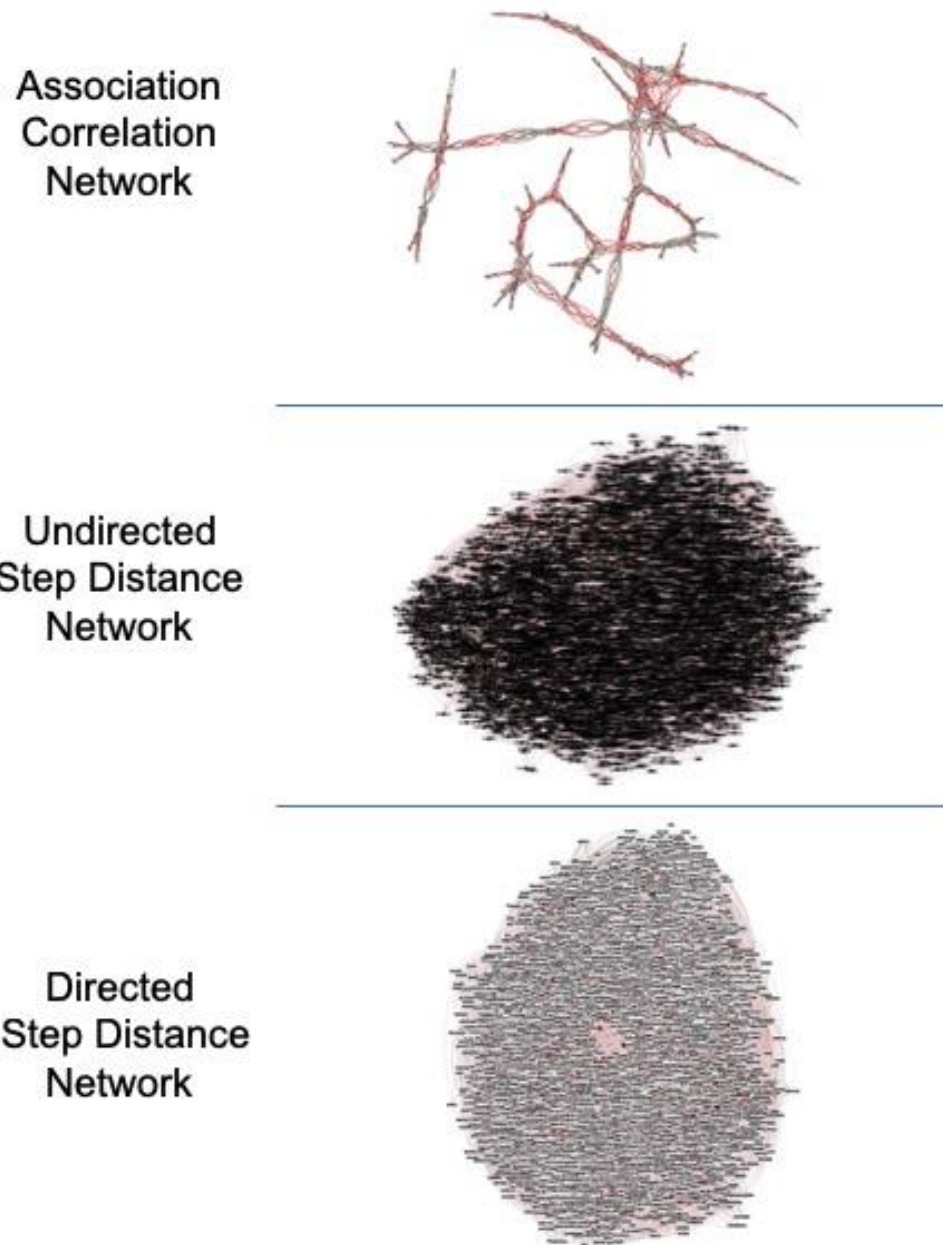


Figure 1. Large-scale visualization of the association-correlation network (top), undirected (middle) and directed (bottom) step distance networks.

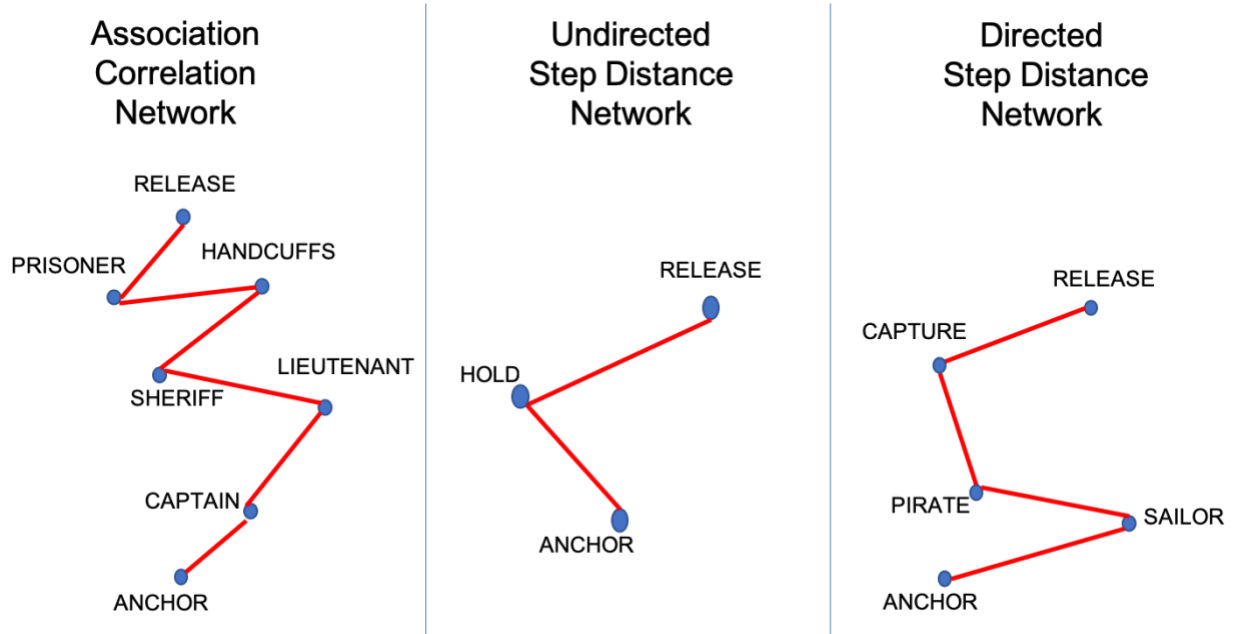


Figure 2. Shortest path from RELEASE to ANCHOR in the association-correlation network (left), undirected (middle) and directed (right) step distance networks.

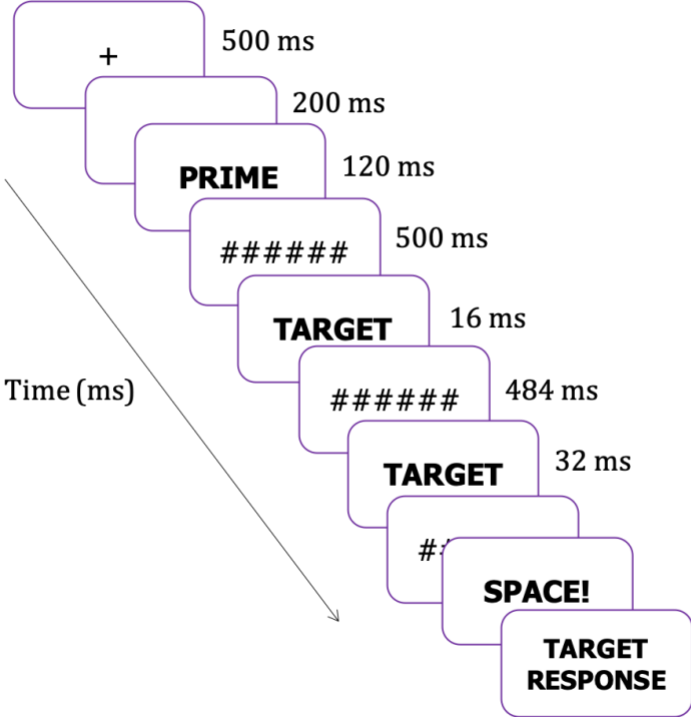


Figure 3. Paradigm for the progressive demasking procedure used in Experiment 2.

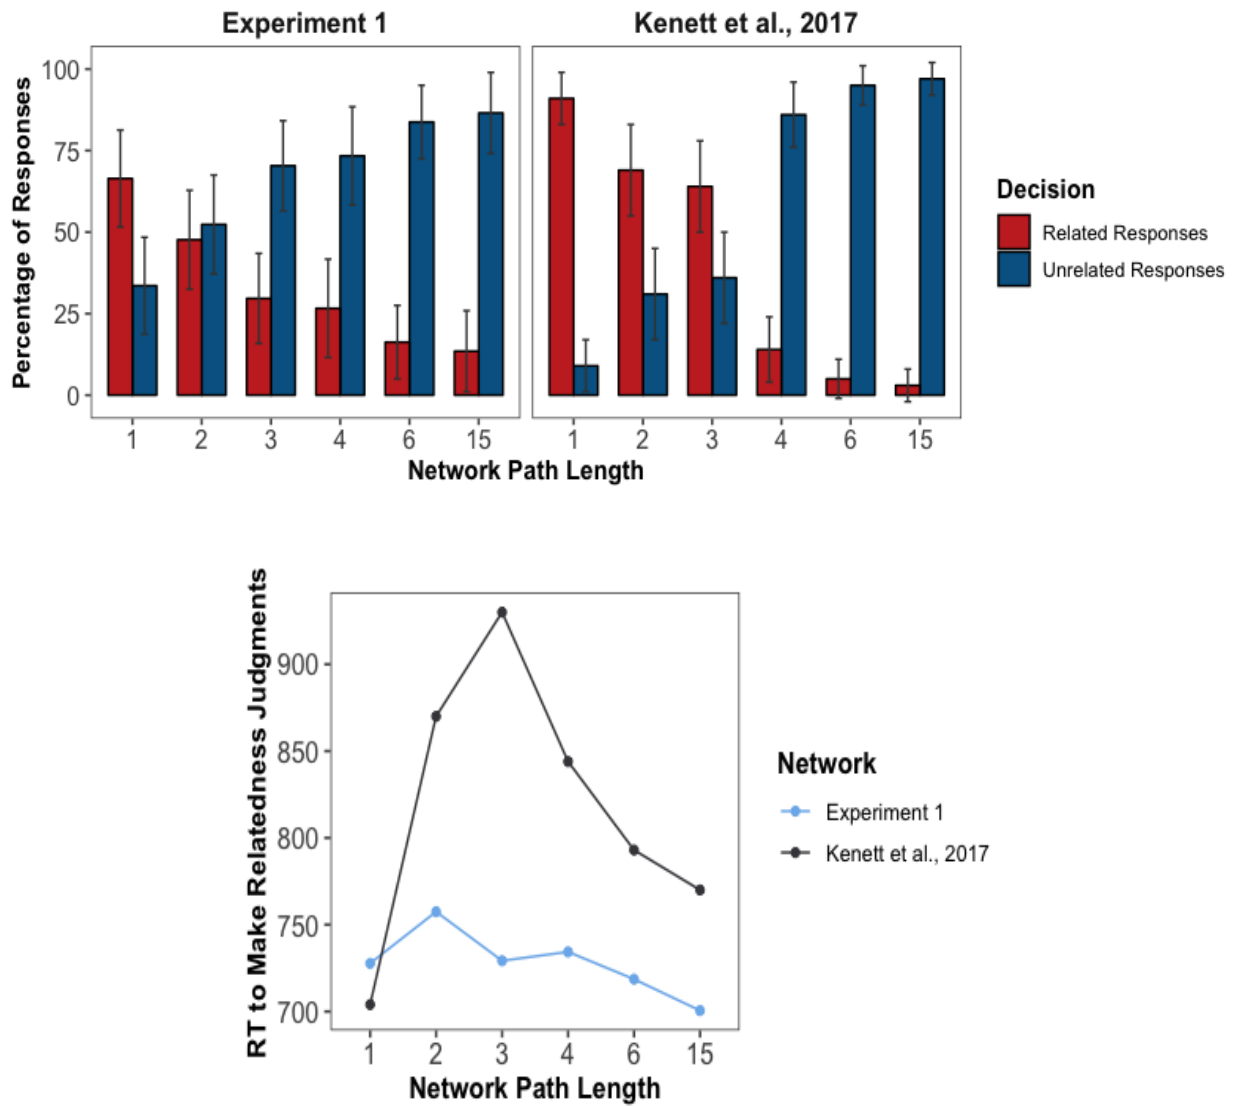


Figure 4. Percentage of related and unrelated responses (Top panel) and response times for relatedness judgments (Bottom panel) in Experiment 1 and Kenett et al. (2017). Error bars represent standard deviations.

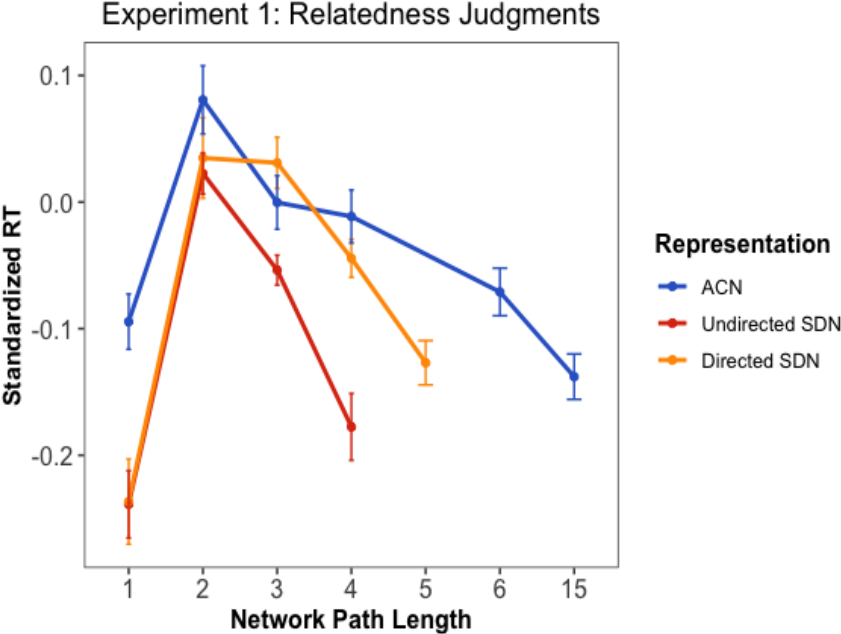


Figure 5. Standardized RTs for relatedness judgments in Experiment 1 as a function of network path lengths. Error bars represent standard errors of the mean.

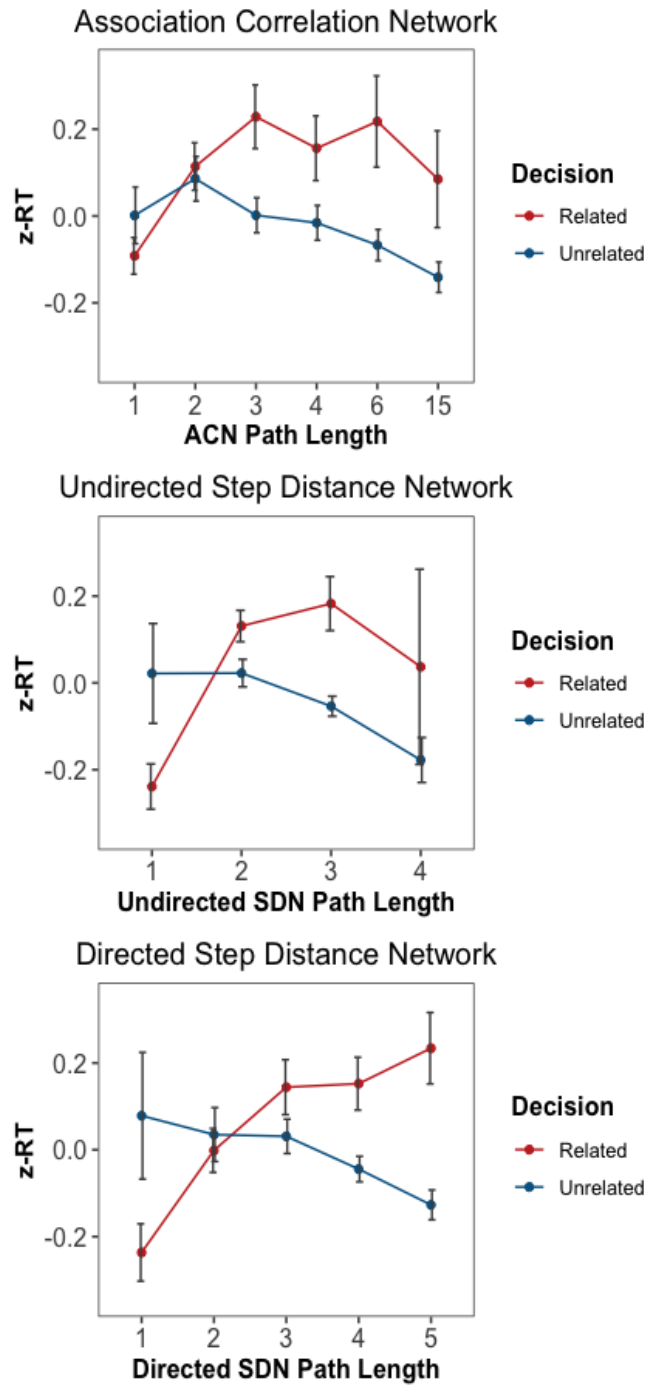


Figure 6. Standardized RTs for relatedness judgments in Experiment 1 as a function of network path lengths and type of relatedness decision. Error bars represent standard errors of the mean.

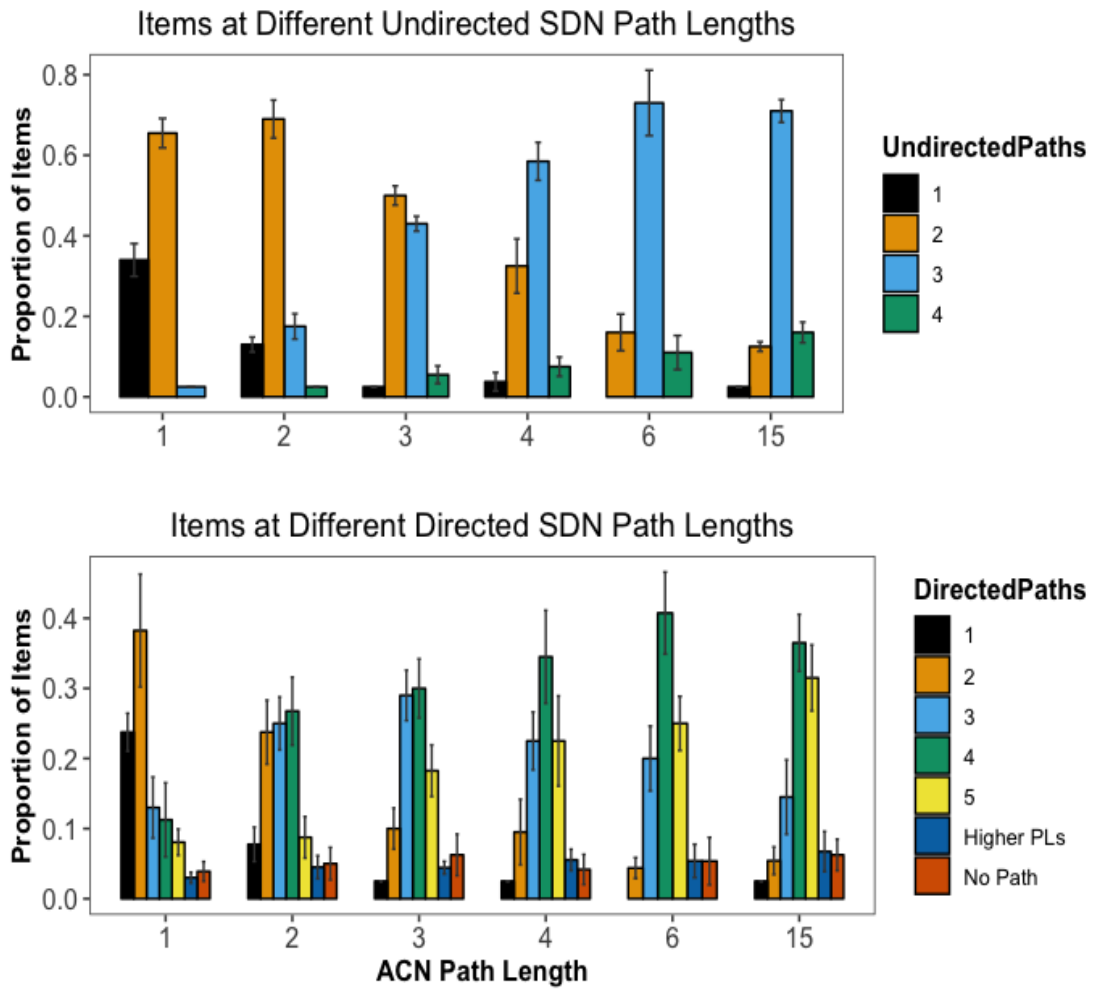


Figure 7. Proportion of items in the different path lengths in the step distance networks (SDNs), as a function of ACN path lengths in Experiment 1. Error bars represent standard errors of the mean.

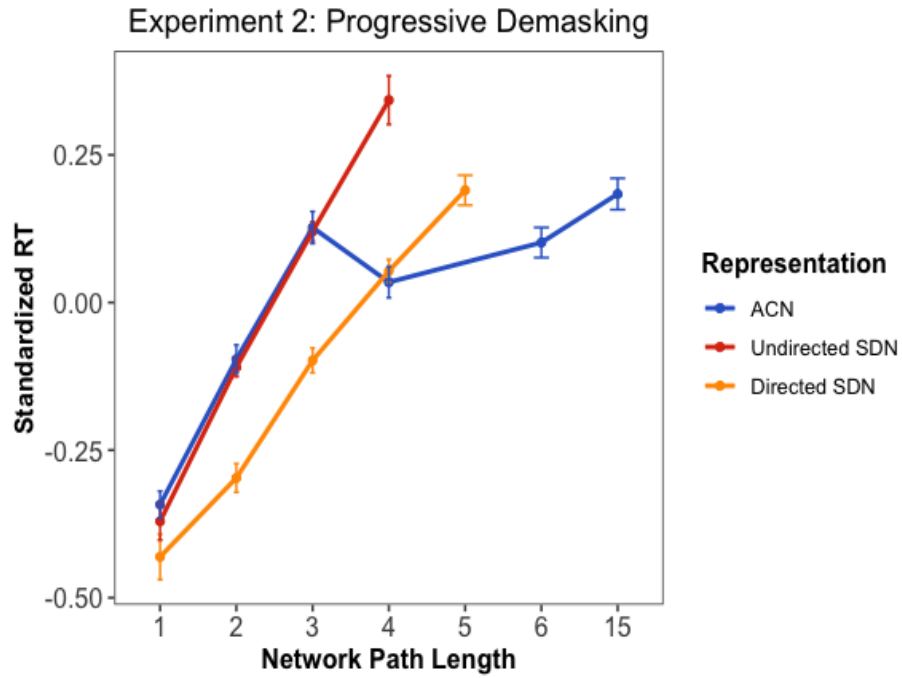


Figure 8. Standardized RTs to identify target word in demasking in Experiment 2 as a function of network path lengths. Error bars represent standard errors of the mean.

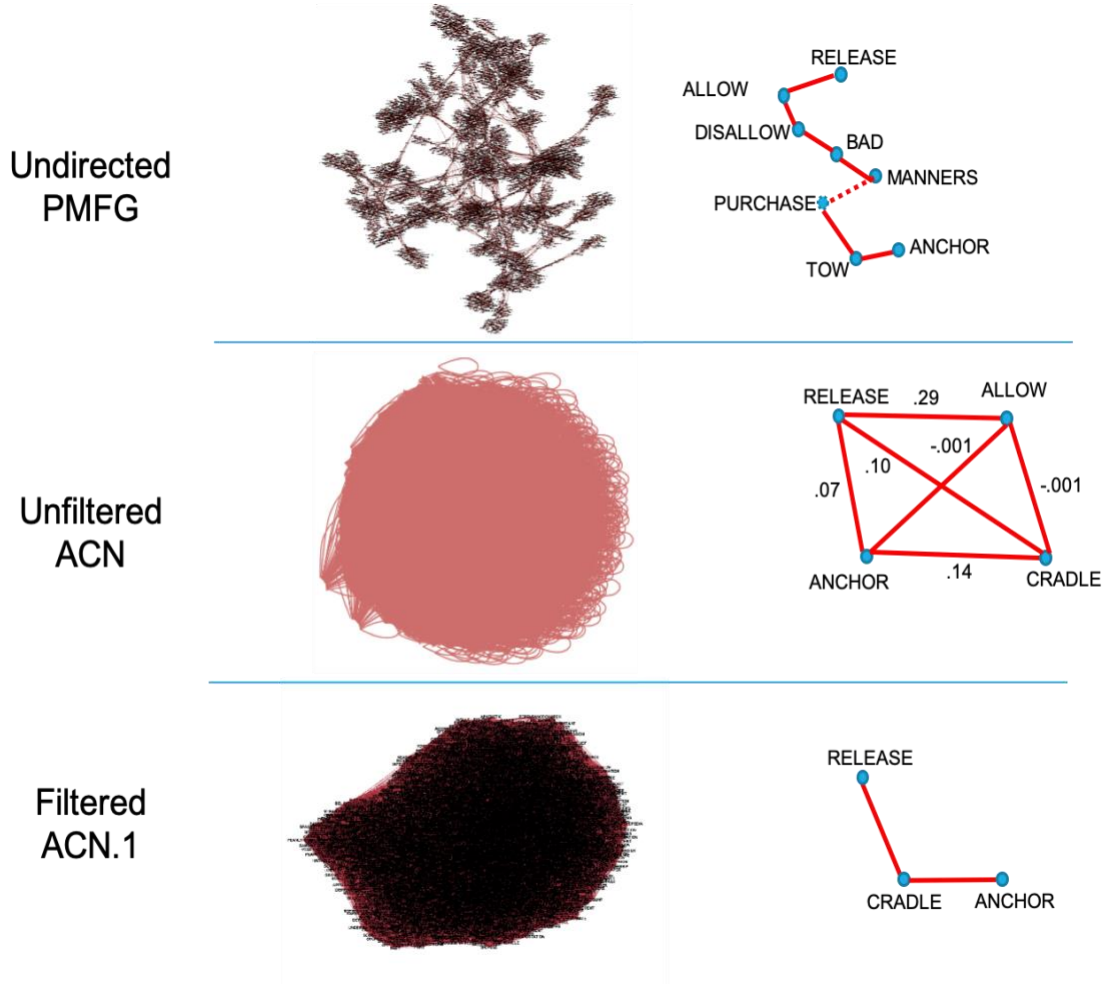


Figure 9. Large-scale visualizations and paths from RELEASE to ANCHOR in the Undirected PMFG (top), Unfiltered ACN (middle) and Filtered ACN.1 (bottom) networks.