

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Distant Concept Connectivity in Network-Based and Spatial Word Representations

Permalink

<https://escholarship.org/uc/item/4tn149vv>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 41(0)

Authors

Kumar, Abhilasha A.

Balota, David A.

Steyvers, Mark

Publication Date

2019

Peer reviewed

Distant Concept Connectivity in Network-Based and Spatial Word Representations

Abhilasha A. Kumar (abhilasha.kumar@wustl.edu)

Department of Psychological & Brain Sciences, Washington University in St Louis, MO 63130 USA

David A. Balota (dbalota@wustl.edu)

Department of Psychological & Brain Sciences, Washington University in St Louis, MO 63130 USA

Mark Steyvers (mark.steyvers@uci.edu)

Department of Cognitive Sciences, University of California, Irvine, Irvine, CA 92697 USA

Abstract

It is presently unclear how localized, word association network representations compare to distributed, spatial representations in representing distant concepts and accounting for priming effects. We compared and contrasted 4 models of representing semantic knowledge (5018-word directed and undirected step distance networks, an association-correlation network and *word2vec* spatial representations) to predict semantic priming performance for distant concepts. In Experiment 1, response latencies for relatedness judgments for word-pairs followed a quadratic relationship with network path lengths and spatial cosines, replicating and extending a pattern recently reported by Kenett, Levi, Anaki, and Faust (2017) for an 800-word Hebrew network. In Experiment 2, response latencies to identify a word through progressive demasking showed a linear trend for path lengths and cosines, suggesting that simple association networks can capture distant semantic relationships. Further analyses indicated that spatial models and correlation networks are less sensitive to direct associations and likely represent more higher-level relationships between words.

Keywords: neural networks; *word2vec*; semantic priming; semantic space model; word association; network science.

Introduction

Understanding language requires the retrieval of meaning from underlying semantic representations of words. A class of models of semantic memory represent words as nodes in a large memory network, where words with similar meanings are connected to each other via edges (see Kenett, Kenett, Ben-Jacob & Faust, 2011; Steyvers & Tenenbaum, 2005). Semantic network models propose localized word representations, in contrast to feature-based or distributed space models (Smith, Shoben & Rips, 1974; Landauer & Dumais, 1997).

Spatial models of semantic memory represent words in a multi-dimensional space, where words are an aggregate of the individual dimensions of the space. The spatial dimensions are derived from statistical co-occurrences in natural language. For example, Latent Semantic Analysis (LSA; Landauer & Dumais, 1997) is a distributional model that measures semantic similarity by calculating co-occurrences of words in a text corpus. LSA successfully simulates complex human behavior in a variety of cognitive tasks but

has had difficulty accounting for semantic priming effects (Hutchison et al., 2008) and power laws (Steyvers & Tenenbaum, 2005), suggesting that spatial models may have some limitations.

A more recent spatial model, *word2vec* (Mikolov, Chen, Corrado & Dean, 2013) has received considerable attention in the fields of computer science and natural language processing. *word2vec* uses neural networks to compute continuous vector representations of words. These semantic representations can then be used to compute an index of semantic similarity between words via vector cosines (higher cosines indicate greater semantic similarity). Interestingly, *word2vec* is able to solve verbal analogy problems (e.g., king: queen::man:?) using simple vector arithmetic, although recent research suggests that *word2vec* successfully captures only certain, simpler types of semantic relationships and not others (Chen, Peterson & Griffiths, 2017). The question of whether individuals use an association-based representation or represent meaning in a high-dimensional space is currently controversial (Griffiths, Steyvers & Tenenbaum, 2007; Jones, Gruenenfelder & Recchia, 2011). Thus, direct comparisons among different types of meaning representations and how they account for more distant semantic relationships is an important next step for the field.

Recently, Kenett, Levi, Anaki and Faust (2017) used a semantic relatedness task to explore the impact of network path length derived from an 800-word Hebrew semantic network. The Hebrew network was created using correlations from continuous free association responses of 60 participants to 800 target words (for complete methodology, see Kenett et al., 2011). The results from the semantic relatedness task indicated that as network path length between word pairs (i.e., shortest distance between two words in the network) increased, fewer word pairs were judged as related. They also reported a quadratic relationship between path length and response latencies to make relatedness judgments, such that response times (RTs) increased for word pairs at shorter path lengths (e.g., BUS-CAR), but after path length 3, RTs systematically decreased for word pairs at longer path lengths (e.g., CHEATER-CARPET). They also showed that this network outperformed LSA and another measure of semantic distance, Positive Pointwise Mutual Information (PPMI) in explaining task performance. However, given that Kenett et

al. used a novel association-correlation methodology based on a Hebrew network, it remains unknown how simpler association networks (e.g., Steyvers & Tenenbaum, 2005) and more recent spatial models (e.g., Mikolov et al., 2013) capture such distant semantic relationships. 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 compare and contrast the structural differences between three different network-based models and the *word2vec* model, across two behavioral tasks. It is important to note here that we do not claim that association-based networks are a complete account of semantic memory, but the issue we are interested in whether networks created from simple associations can indeed capture distant semantic priming effects, and how they compare to other models of semantic memory, such as spatial models and the association-correlation network. There is a rich tradition of using network-based models to accommodate priming effects (Anderson, 2000; Collins & Loftus, 1975), and we were mainly interested in comparing different types of network-based approaches to each other in accounting for this well-studied task, and also to other spatial representations. In Experiment 1, we extended and replicated the patterns reported by Kenett et al. in the Hebrew semantic relatedness task in three large semantic networks in English along with cosines from the *word2vec* model. We created these networks from a 5018-word database of free association norms collected by Nelson, McEvoy & Schreiber (2004) to examine the extent to which network path lengths would predict performance in the relatedness judgment task.

A potential concern regarding the performance of network models created through human association norms in Experiment 1 is both relatedness judgments and word associations direct attention to the meaning dimension, and thus the patterns observed may just be due to overlap in the type of task. Further, the quadratic pattern observed may just reflect how the semantic “distance” between two words might influence the related/unrelated decision and how a particular individual partitions items into these arbitrary categories. We attempted to address this concern by employing a task that does not require accessing meaning-related information to make the response. Thus, in Experiment 2, participants first viewed a briefly presented prime (120 ms) and then identified targets through a visual demasking task. Hence, we were able to directly compare the different network configurations and spatial representations in accounting for performance in two behavioral tasks.

Semantic Network Construction

To construct the semantic networks, we used a 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. The cues were selected by Nelson et al. after multiple rounds of data collection, and typically, the most

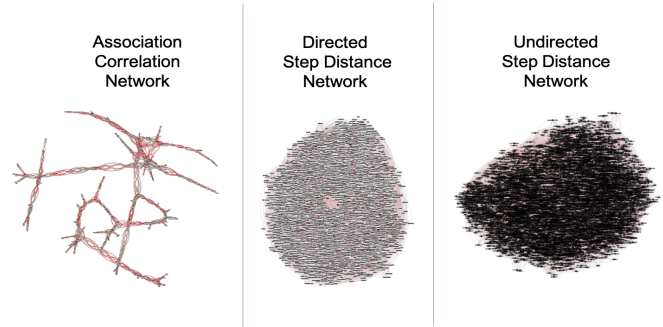


Figure 1: Large-scale visualization of the Association-Correlation Network, Directed and Undirected Step Distance Networks.

frequent responses for each of the cues were contained within the 5018 cues themselves. Responses were included only if at least two participants produced the same response, thus excluding idiosyncratic responses from the database. Responses that were not within the 5018 cues were also excluded during network construction. We constructed three networks from this database: an Association-Correlation Network (ACN), an Undirected Step Distance Network (Undirected SDN) and a Directed Step Distance Network (Directed SDN).

Association-Correlation Network

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 word. 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, resulting in a smaller planar network (a network in which no edges cross each other) with the same target nodes, but only edges that represent the most relevant associations between target words. Path length between word-pairs was then calculated as the shortest path from one word to another in this smaller network. Figure 1 (Left 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

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. In the Undirected SDN, words were connected if *a* evoked *b* or *b*

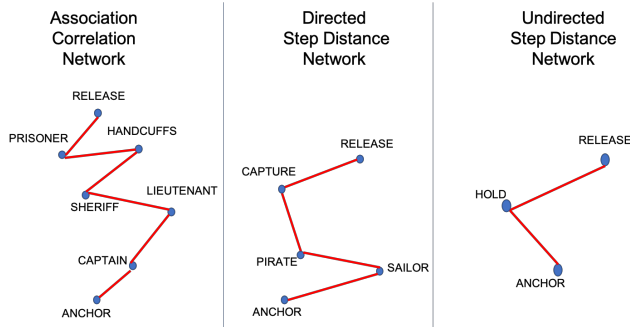


Figure 2: Shortest path from RELEASE to ANCHOR in the Association-Correlation Network, Undirected and Directed Step Distance Networks.

evoked a , independent of the associative direction. Path length for each word pair in the network was calculated as the shortest path from one word to another. Figures 1 and 2 (Middle and right panels) display visualizations of the two SDNs, and the shortest path from RELEASE to ANCHOR.

Network Comparisons

Table 1 displays the network parameters for the three networks. As is evident from the large-scale visualizations, 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 distant 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, there were considerable differences across the different types of network configurations. As shown in Figure 1, the ACN is a sparsely connected network, in which obscure,

Table 1: Network parameters for the semantic networks

	Simple Step Distance Networks		Association-Correlation Networks	
	Undirected	Directed	English	Hebrew
n	5018	5018	5018	800
$\langle k \rangle$	22	12.7	5.85	5.94
L	3.04	4.27	23	10
D	5	10	61	25
C	.186	.186	.69	.68
L_{random}	3.03	4.26	1.95	3.94
C_{random}	.004	.004	.05	.005

Note. n = the number of nodes; $\langle k \rangle$ = average number of connections; L = average shortest path length; D = diameter of network; C = clustering coefficient; L_{random} = 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 for network path lengths and $word2vec$ cosines for word-pairs in Experiments 1 and 2

	ACN	Undirected SDN	Directed SDN	$word2vec$ Cosines
ACN	1	-	-	-
Undirected SDN	.49	1	-	-
Directed SDN	.35	.58	1	-
$word2vec$ Cosines	-.42	-.55	-.45	1

Note: All correlations were significant at the $p < .05$ level

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) are exaggerated. Overall, path lengths derived from the two SDNs were very highly correlated, suggesting that the simple associative networks largely overlap in their network structure, and differ from the ACN.

Vector Cosines via $word2vec$

The $word2vec$ model (Mikolov et al., 2013) trains neural networks based on words that naturally co-occur in a text corpus and uses this contextual information to predict a word’s immediate contextual neighborhood. Typically, these contextual words have probabilities associated with them, which indicate the likelihood of words co-occurring together in natural language. If two words occur in similar contexts, the model learns similar vector representations for those words. Cosines between these vector representations thus serve as indices of semantic similarity. For all the word pairs used in the current experiments, we obtained $word2vec$ cosines from a pre-trained model trained on 100 billion words from a Google News dataset (Mikolov et al., 2013). Table 2 displays the correlations between $word2vec$ cosines and path lengths derived from the three networks described above. Note that $word2vec$ cosines were negatively correlated with the path lengths, due to the direct cosine similarity measure used. Further, there were considerable differences across the models in the extent to which they captured “semantic similarity”, given that the average correlation among all the measures was only .46.

Experiment 1

Methods

Participants Forty Amazon Mechanical Turk users ($M_{age}=37$ years, $SD = 10.4$) and 40 undergraduate students ($M_{age}=20$ years, $SD = 0.8$) recruited from Washington University in St Louis participated in the study. All participants were self-reported native English speakers.

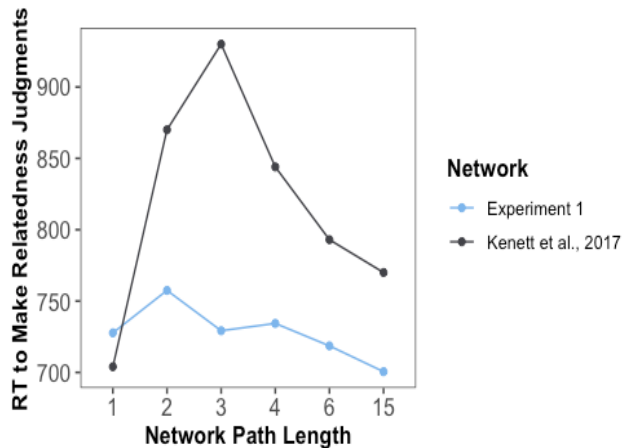


Figure 3: Response times for relatedness judgments in Experiment 1 and Kenett et al. (2017)

Materials In order to extend and replicate the Kenett et al. study, we randomly sampled 40 word-pairs from path lengths 1, 2, 3, 4, 6 and 15 from the ACN. The stimuli consisted of 1200 distinct word-pairs across 5 lists. For each word-pair sampled from the ACN, we also obtained path lengths in the Undirected and Directed SDN and obtained vector cosines from the *word2vec* model. We also obtained lexical characteristics (word length, frequency, lexical decision times and concreteness) for all the words from the English Lexicon Project (ELP; Balota et al., 2007) and used these as covariates in our analyses. All items used in the current study are available at <https://github.com/abhilasha-kumar/Distant-Semantic-Connectivity>.

Procedure

The relatedness task was developed in JSPsych, an online software for conducting psychological experiments. Each participant completed the experiment online. Following Kenett et al., on each trial, participants saw a fixation cross for 200 ms, followed by a blank screen for 100 ms. Then, the prime was briefly presented for 120 ms, followed by the target for 120 ms. Participants decided whether the prime and target were related or unrelated and responded by pressing a key (K or L, counterbalanced). After a response, participants saw a blank screen for 500 ms before the next trial.

Results

There were no differences in the overall patterns between the five lists, or the Amazon Mechanical Turk or Washington University sample, thus all analyses included the full sample.

Effect of ACN Path Length on RTs To replicate the analysis procedures reported in Kenett et al. (2017), each path length was first classified as related or unrelated, based on the percentage of related and unrelated responses to specific word pairs. The following were the percentages of “related” responses for the path lengths: 1 (66%), 2 (47%), 3 (29%), 4 (27%), 6 (16%) and 15 (13%). Based on these percentages

and the criterion of at least 50% of words producing a related response, only path length 1 was considered related, and the remaining path lengths were considered unrelated. To minimize any effects of slowing and individual differences, all 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. A repeated measures Analysis of Variance (ANOVA) on mean RT revealed a significant main effect of path length, $F_1(5, 395) = 7.42, p < .001, \eta_p^2 = .09$. RTs significantly increased from path length 1 to 2 ($p = .006$), decreased from path lengths 2 to 3 ($p = .001$) and 4 to 15 ($p = .015$). As shown in Figures 3 and 4, we successfully replicated the pattern reported by Kenett et al. for RTs as a function of path length in the ACN. Importantly, this pattern persisted after including degree of relatedness as a predictor in our analyses, standardizing the RTs and controlling for lexical variables such as word frequency, length, concreteness and standardized lexical decision times, as well as mean degree (i.e., number of direct neighbors of the words) using linear mixed effects models.

Effect of SDN Path Length on RTs In addition to the ACN based on Kenett et al., as noted, we also examined the effect of path lengths derived from two SDNs (Undirected and Directed) based on the method used in Steyvers and Tenenbaum (2005) on standardized RTs in the relatedness task. As shown in Figure 4, both the Undirected and Directed networks also showed a quadratic trend for standardized RTs as a function of path length, with RTs significantly rising from path lengths 1 to 2 ($p < .001$) and then reliably decreasing

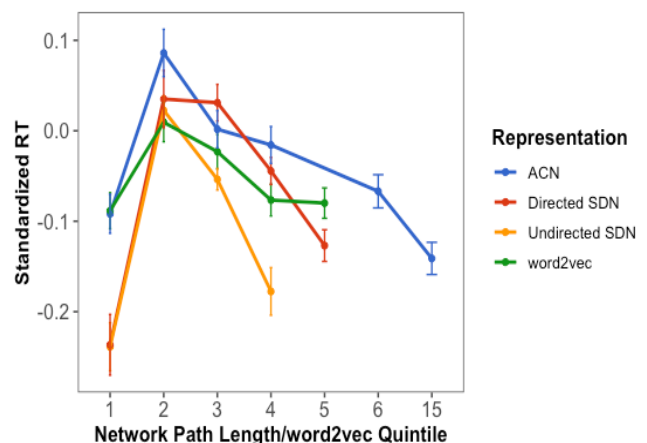


Figure 4: Standardized RTs for relatedness judgments in Experiment 1 as a function of network path lengths and *word2vec* cosine quintiles (reverse-scored)

from path length 2 onwards. We observed a significant decline in RTs from path lengths 3 to 4 in the Undirected ($p < .001$), and from 2 to 5 in the Directed network ($p < .001$).

Effect of word2vec Cosines on RTs We also computed vector cosines derived via *word2vec* for each of the word pairs in Experiment 1. As shown in Figure 4, continuous *word2vec* cosines successfully predicted standardized RTs to make relatedness judgments ($b = -.22, t = -3.54, p < .001$), and reproduced the quadratic pattern previously observed.

Discussion

The results from Experiment 1 provide strong evidence for multiple-step priming in the relatedness judgment task, and also replicate and extend the quadratic pattern observed by Kenett et al. (2017) for the Hebrew network. In addition, simple directional and nondirectional SDNs also captured distant semantic relationships between concepts. This is noteworthy, as it indicates that the number of “steps” in the ACN do not necessarily reflect *direct* associative strength, at least based on distances captured by simple SDNs. Of course, this does not imply that the ACN distances are unimportant, as the ACN shows comparable performance in the current task. We also found that the *word2vec* model successfully captured the quadratic trend, although there do seem to be differences in the semantic information captured by all the models, based on the relatively low correlations across the networks.

It is important to note that the nature of the relatedness decisions is likely driving the quadratic trend. Specifically, RTs are slowed to make “unrelated” decisions for the more ambiguous items e.g., at path lengths 2 and 3. Interestingly, the RTs for only the “related” decisions continued to increase with greater path lengths, a finding that is more consistent with a spreading-activation account. In addition, the networks in this study were explicitly created from free association norms, and their explanatory power may reflect the high degree of overlap between the base task (free association) and the relatedness judgment task. Thus, in Experiment 2, we explored whether network path length and vector cosines can account for semantic priming in a primed progressive demasking task, which does not explicitly involve explicit semantic retrieval to make a response.

Experiment 2

Methods

Participants Thirty-nine young adults ($M_{age} = 20.9$ years, $SD = 2.8$) were recruited from undergraduate courses at Washington University in St Louis. All participants were Native English speakers.

Materials One list of 240 items was randomly chosen from one of the five lists used in Experiment 1. As before, the list contained 40 word-pairs from path lengths 1, 2, 3, 4, 6 and 15 from ACN. Each word pair also had corresponding path lengths in the undirected and Directed SDN, as well as

word2vec cosines. This list was then used to create two lists counterbalanced across participants, so that each word was a prime as well as a target in the study.

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, displayed for 120 ms. Immediately after, the target word was progressively demasked on the screen. During progressive demasking, the display alternated between the target (e.g., XXXX) 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 target display time progressively increased. 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 and typing in the target word. The next trial began immediately after typing the target and pressing spacebar.

Results

Effect of ACN Path Length on RTs All trials in which the correct target was not identified were excluded from analyses (2.7%). Next, we standardized the RTs to identify the target as in Experiment 1. A repeated measures ANOVA revealed a significant effect of path length, $F(5,190) = 53.85, p < 0.001, \eta_p^2 = .586$. As shown in Figure 5, we observed a significant increase in RTs from path lengths 1 to 2 ($p < .001$), and 2 to 3 ($p < .001$). Differences between RTs at path length 3 and higher ACN path lengths were not reliable. These effects persisted after controlling for lexical variables & mean degree (i.e., number of direct neighbors of the words).

Effect of SDN Path Length on RTs. We also examined the effect of path lengths from the Undirected and Directed SDNs on standardized RTs. As shown in Figure 5, path lengths from the Undirected SDN significantly predicted RTs to identify the target. RTs increased from path length 1 to 2 ($p = .001$), from path lengths 2 to 3 ($p < .001$), and then marginally from 3 to 4 ($p = .058$). Path lengths from the Directed SDN also predicted RTs to identify the target. RTs increased from path lengths 2 to 3 ($p = .015$) and 4 to 5 ($p = .038$).

Effect of word2vec cosines on RTs We also obtained vector cosines derived via *word2vec* for each of the word pairs, as in Experiment 1. As shown in Figure 5, continuous *word2vec* cosines also successfully predicted standardized RTs to identify the target ($b = -1.34, t = -9.18, p < .001$).

Model Comparisons Because the results from this task were not complicated by the relatedness decision as in Experiment 1 (i.e., RTs should be linearly related to demasking performance), we were able to directly compare the model

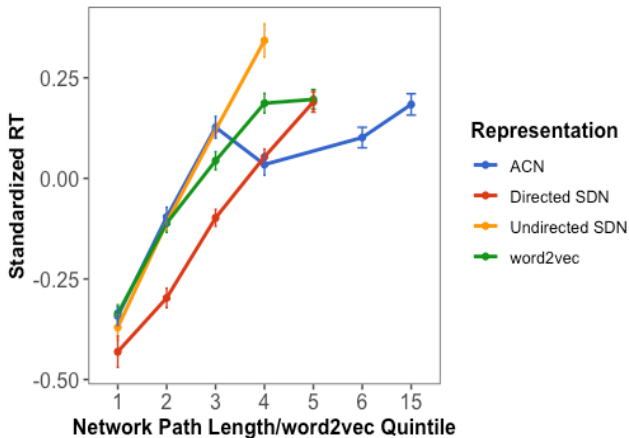


Figure 5: Standardized RTs to identify target word in demasking in Experiment 2 as a function of network path lengths and *word2vec* cosine quintiles (reverse-scored)

estimates. To estimate the unique variance accounted for by each type of network configuration at the item level, we calculated the individual R^2 for each model, as well as estimates of AIC and BIC, after controlling for covariates. As shown in Table 3, the models had overall comparable fits, and explained a significant amount of variance over and above the model with just the covariates, although as discussed before, these models seem to capture somewhat different semantic information.

Discussion

Results from Experiment 2 indicated that network path lengths can indeed account for performance in a progressive demasking task. RTs linearly increased as a function of SDN path lengths and *word2vec* cosines. This is especially interesting as the demasking task does not require any direct retrieval of semantic association to make the response, and yet, we see that path lengths derived from word associations directly predict demasking response latencies. Further, we found reliable differences at relatively distant path lengths in the simple association networks, suggesting that simple association networks are able to capture distant semantic relationships in the memory network, even in tasks that do not necessarily direct attention to semantics. Interestingly, we find that path lengths from the ACN increase linearly only up to 3 steps, after which the network seems to no longer be sensitive to priming effects in this task, suggesting differences in the network structures.

General Discussion

A primary goal of the present study was to compare the extent to which measures of semantic similarity derived from different types of network-based models explained distant semantic priming. In Experiment 1, we replicated and extended a pattern previously reported by Kenett et al. (2017) to a larger 5018-word association network in English and also

Table 3: Model comparison metrics for Experiment 2

Model	R^2 (%)	AIC	BIC	Likelihood ratio test
Covariates	13.33	561.9	586.1	---
ACN	26.99	500.8	545.1	$p < .001$
U-SDN	22.16	523.3	559.6	$p < .001$
D-SDN	25.98	506.4	550.8	$p < .001$
<i>word2vec</i>	28.03	486.8	515	$p < .001$

compared their graph-theoretical approach of capturing semantic similarity with simpler Undirected and Directed Step Distance Networks (Steyvers & Tenenbaum, 2005). Our results indicated that simple association networks can also capture similar distant relationships between words in the lexicon. Experiment 2 indicated that network models also successfully capture performance in tasks that do not directly rely on word association.

As described earlier, the ACN uses correlations between association responses and a planarity criterion to construct the network, and possibly captures more higher-level associations. This leads to several direct word associations (e.g., TIGER-STRIPES is 37 steps away in the ACN and 1 step away in the SDNs) being dropped, giving rise to more high-level associations (e.g., SUEDE-SERPENT is only 2 steps away in the ACN but farthest, i.e., 4 steps away in the SDNs). The SDNs, on the other hand, capture *direct* associations between words. Importantly, given that all networks had comparable fits, it seems that each network captured different sources of variance in the task.

It is possible that the ACN may be differentially sensitive to semantic relationships if a different criterion for network construction was used, or possibly in a conceptually driven semantic task, which would suggest that different types of stimuli/tasks emphasize different properties of the lexicon. Indeed, Gruenfelder, Recchia, Rubin and Jones (2015) recently argued for a hybrid representation of 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 semantic representations in tasks that do not explicitly involve word association but are still sensitive to semantic relationships.

Another important goal of the current study was to investigate how network-based models of semantic representation compare to a distributed model, *word2vec*, which has been shown to explain human performance in some semantic tasks. Our results indicate that *word2vec* successfully captures similar patterns of behavior as the semantic networks. However, we also observed important differences in the semantic relationships captured by each of the models. For example, the word BOXING is 2 steps away from the word SPLINTER in the Undirected SDN but is very weakly associated in the *word2vec* space with a cosine of -0.022. Thus, there appear to be differences in the type of semantic information the models capture, e.g., the path from

BOXING to SPLINTER is mediated by the word PAIN in the association networks, but it is possible that this particular usage of SPLINTER does not co-occur in the same contexts as BOXING very often, which is the mechanism underlying *word2vec* model. Thus, these findings indicate that the nature of the task as well as the underlying representation are both critical variables that determine the extent to which semantic models explain human performance. Importantly, the tasks in the current study focused on semantic priming, and it is possible that spatial models and correlation networks are most useful in conceptual tasks like verbal analogies.

There were some limitations to the current study. First, the Hebrew network used in Kenett et al. (2017) was based on responses from a continuous free association task, whereas the Nelson et al. norms are based on a discrete free association task. The validity of both continuous and discrete responses has been debated (Hahn, 2008; Nelson, McEvoy & Dennis, 2000) and our use of discrete responses may have produced a different network structure than one based on continuous responses. However, given that the English ACN and SDNs were created from the same norms, we believe that the differences observed between the ACN and the SDNs were not critically influenced by the nature of associative responses per se, although this issue deserves further exploration. Further, the *word2vec* model was trained on a Google News corpus, which is very different from the Nelson et al. database, and the type of corpus can impact how well semantic models account for performance (Recchia & Jones, 2009). Thus, the nature of the task, stimuli and training corpora are all likely to influence the extent to which semantic models explain cognitive task performance.

In conclusion, the current set of experiments investigated the predictive power of path lengths derived from three large semantic networks and cosines derived from a neural network model in two behavioral tasks and provided strong evidence for multiple-step priming. We also demonstrated important structural differences between correlation-based networks and simple association networks and showed that simple association networks are also able to capture relatively distant semantic relationships. Finally, we showed that *word2vec* successfully captures similar behavioral patterns across two tasks. However, based on preliminary analyses, it appears that *word2vec* and the ACN are more likely to capture higher-level semantic representations, whereas simple step networks are more likely to capture direct associations. Clearly, further work is needed to substantiate these observations.

References

- Anderson, J. R. (2000). *Learning and memory: An integrated approach*. (2nd ed.). New York: Wiley.
- 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.
- 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.
- Griffiths, T. L., Steyvers, M., & Tenenbaum, J. B. (2007). Topics in semantic representation. *Psychological Review, 114*(2), 211.
- 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.
- Jones, M., Gruenenfelder, T., & Recchia, G. (2011, January). In defense of spatial models of lexical semantics. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 33, No. 33).
- 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.
- 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.
- 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.
- Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- 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.
- Recchia, G., & Jones, M. N. (2009). More data trumps smarter algorithms: Comparing pointwise mutual information with latent semantic analysis. *Behavior Research Methods, 41*(3), 647-656.
- 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.