TOPICS
TOPICS IN COGNITIVE SCIENCE

Topics in Cognitive Science 0 (2021) 1–24

ISSN: 1756-8765 online
DOI: 10.1111/tops.12548

This article is part of the topic "Networks of the Mind: How Can Network Science Elucidate Our Understanding of Cognition?" Thomas T. Hills and Yoed Kenett (Topic Editors). For a full listing of topic papers, see http://onlinelibrary.wiley.com/journal/10.1111/(ISSN)1756-8765/earlyview.

A Critical Review of Network-Based and Distributional Approaches to Semantic Memory Structure and Processes

Abhilasha A. Kumar,[a] Mark Steyvers,[b] David A. Balota[a]

[a]*Psychological & Brain Sciences, Washington University in St. Louis*
[b]*Department of Cognitive Sciences, University of California, Irvine*

Received 6 January 2021; received in revised form 14 May 2021; accepted 19 May 2021

Abstract

Some of the earliest work on understanding how concepts are organized in memory used a network-based approach, where words or concepts are represented as nodes, and relationships between words are represented by links between nodes. Over the past two decades, advances in network science and graph theoretical methods have led to the development of computational semantic networks. This review provides a modern perspective on how computational semantic networks have proven to be useful tools to investigate the *structure* of semantic memory as well as search and retrieval *processes* within semantic memory, to ultimately model performance in a wide variety of cognitive tasks. Regarding representation, the review focuses on the distinctions and similarities between network-based (based on behavioral norms) approaches and more recent distributional (based on natural language corpora) semantic models, and the potential overlap between the two approaches. Capturing the type of relation between concepts appears to be particularly important in this modeling endeavor. Regarding processes, the review focuses on random walk models and the degree to which retrieval processes demand attention in pursuit of given task goals, which dovetails with the type of relation retrieved during tasks. Ultimately, this review provides a critical assessment of how the network perspective can be reconciled with distributional and machine-learning-based perspectives to meaning representation, and describes

Correspondence should be sent to Abhilasha A. Kumar, Washington University in St. Louis, CB 1125, One Brookings Drive, St. Louis, MO 63130-4899, USA. E-mail: abhilasha.kumar@wustl.edu

how cognitive network science provides a useful conceptual toolkit to probe both the structure and retrieval processes within semantic memory.

Keywords: Cognitive network science; Distributional semantic models; Semantic memory; Semantic networks

---

Within the past decade, there has been an explosion of research aimed at better understanding semantic memory representations using tools from network science (for a recent review, see Siew, Wulff, Beckage, & Kenett, 2019). The majority of the work in this domain, broadly termed as "cognitive network science," has been directed at understanding how close or far apart are concepts within a network-like model, and ultimately using these measures of *distance* to predict various aspects of human behavior. In addition, there has been considerable interest in using computational metrics from graph theory to describe the local and global characteristics of networks, that is, how concepts are organized and clustered within a network configuration. Understanding *semantic* memory organization and processes (how word meanings are structured within memory and retrieved) has been at the forefront of these explorations. Indeed, the notion of a semantic network has been computationally implemented using behavioral norms from semantic tasks or natural language corpora, giving rise to quantitative estimates of concept connectivity (e.g., Steyvers & Tenenbaum, 2005) and how information or activation may "spread" within a network (e.g., De Deyne, Verheyen, & Storms, 2016; Vitevitch et al., 2011). This burgeoning field of semantic network research provides a new lens to reexamine questions pertaining to both the *structure* of knowledge representation as well as *processes* that operate upon this structure to ultimately produce complex human behavior.

1. Early debates in semantic memory research

Although there has been considerable recent progress in developing computational semantic networks, the notion of a semantic network reflecting the relations among words is not new to cognitive science. A semantic network was first explicitly implemented by Quillian (1967), who attempted to maximize storage space within a computer-like architecture, by conceptualizing a hierarchical network with concepts and edges describing related propositions (i.e., a *robin* <is a> *bird* and a *bird* <has> *wings*). Although there was initial empirical support for such a hierarchical model by Collins and Quillian (1969), there were some limitations that were addressed in the Collins and Loftus (1975) model, which attempted to capture the strength of the relationship between concepts by the length of the pathway connecting two nodes. The early concerns also motivated an important alternative representational format that relied on the notion of a set of semantic *features* representing concepts (e.g., a *bird* <has wings>, <flies>, <lives in a nest>, etc.), instead of an unanalyzable localist node within a network (e.g., Smith, Shoben, & Rips, 1974). Of course, a feature-based representation need not be necessarily orthogonal to a network-based approach, given that a semantic network can also encode features (e.g., *bird-wings*). The critical difference between these representational formats was that feature-based models solely emphasized the shared features between concepts, whereas semantic network connections could, in principle, reflect a broad range of semantic relationships, some of which may not directly emerge out of shared features (e.g., *baby-stork*, *mouse-cheese*, etc.).

An important insight that emerged from these early theoretical debates on the nature of semantic representations (i.e., network vs. feature-based) was that these representational distinctions demanded distinct retrieval *processes*. Collins and Loftus (1975) relied heavily on the metaphor of a *spreading activation mechanism*, that is, concepts within the network were activated and activation then spread to neighboring nodes and propositions in proportion to the strength of association between them, leading the individual to confirm or reject the validity of the sentence in a sentence verification task. On the other hand, Smith et al. proposed that individuals engaged in a *feature comparison process* between the subject and predicate of the sentence, which was followed by a two-stage decision process that involved both analytic (more attention-based) and nonanalytic (more automatic) processes used to drive responses (see McCloskey & Glucksberg, 1979 for an early single process feature comparison model). These original explorations into semantic *representations* nicely illustrate how it is critical to consider the nature of the representation (e.g., network vs. feature-based) along with the specific *processes* that access the representation (e.g., spreading activation vs. featural comparison) to account for behavior.[1]

The present review will examine how modern computational network models of semantic memory have advanced our understanding of how concepts are represented and retrieved within semantic memory. The first section will provide an overview of some basic representational principles of recent developments in network science, emphasizing the differences and similarities between network representations and more recent distributional semantic representations. The second and third sections will delineate how semantic network research has provided useful tools to investigate both the *structure* of semantic memory as well as the *processes*, respectively, which operate upon this structure to ultimately provide an account of cognitive behavior across a variety of tasks. The fourth section will discuss how one might disentangle data, structure, and processes by providing a brief overview of the strengths and limitations of the network-based approach in comparison to the distributional modeling approach. The final section will address some conceptual questions about networks and their utility, to ultimately provide ways in which network-based perspectives can help develop a more unified account of semantic memory structure and processing.

2. Network versus distributional models of meaning representation

An important class of models in addition to feature-based and network-based models are distributional semantic models (DSMs). These include a family of computational models that apply statistical algorithms to word co-occurrence patterns extracted from large text corpora (e.g., Wikipedia or Google News databases) to *learn* semantic representations. Although one could argue that DSMs date back to early work by Osgood (1952), there has been an

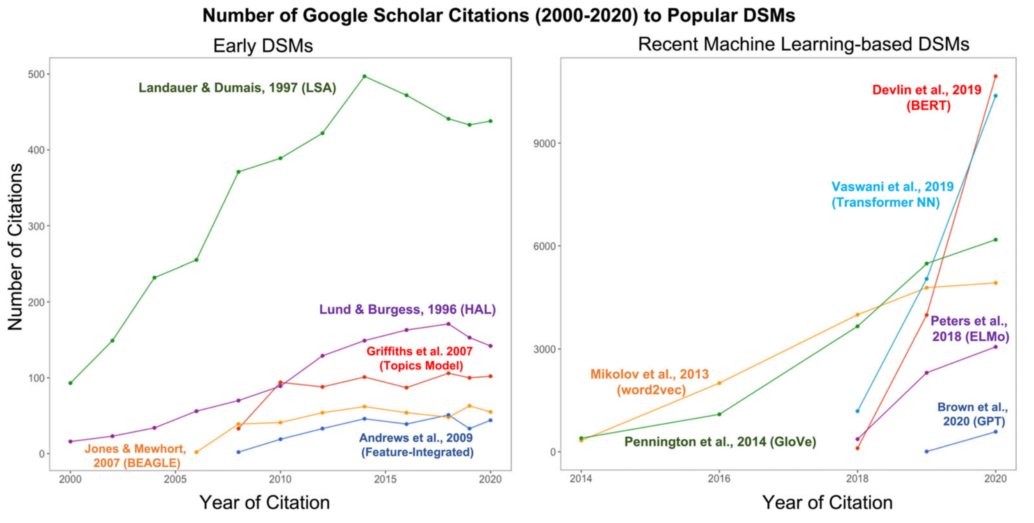

Fig. 1. Number of citations based on Google Scholar to popular Distributional Semantic Models (DSMs) from 2000 to 2020.

explosion of these models within the past decade due to the availability of large text databases and the implementation of different types of machine learning algorithms. As shown in Fig. 1, the left panel indicates some of the first DSMs that were developed to capture high-dimensional semantic spaces based on natural language corpora. The right panel of Fig. 1 reflects more recent DSMs that are based on larger natural language corpora and involve machine learning algorithms to generate the semantic space. It is particularly noteworthy that there is an explosion of citations to these more recently developed models within the past decade (note the different scales for the y-axes across the two panels).

The guiding principle within DSMs is that meaning can be inferred via statistical regularities in natural language, which can be captured through different learning mechanisms such as Hebbian or error-driven learning, probabilistic inference, etc. (more recent work also examines such regularities in other modalities such as vision, see Bruni, Tran, & Baroni, 2014). Within DSMs, meaning is typically represented in a high-dimensional space (e.g., BEAGLE; Jones & Mewhort, 2007), or a large collection of topics (as in topic models, see Griffiths, Steyvers, & Tenenbaum, 2007). Relationships between concepts are then inferred via some measure of association strength (e.g., cosine similarity between semantic vectors, probability distributions within topic models, etc.) and DSMs have shown remarkable success at modeling some aspects of human behavior (for recent reviews, see Günther, Rinaldi, & Marelli, 2019; Kumar, 2021b).[2]

Clearly, the notion of extracting patterns of co-occurrences from natural language in DSMs is very different from a feature-based and network-based perspective. It is important to emphasize here these classes of models (network-based, feature-based, and distributional) generally vary along two critical dimensions—the underlying *data* source, as well as the ultimate *structure* of the representation. In terms of *data*, while distributional models typically use natural text or speech corpora to *learn* semantic representations, network and feature-based models are typically based on human-generated norms. The word "typically" is important here because one can generate semantic networks from text corpora (e.g., Beckage, Smith, & Hills, 2011; Steyvers & Tenenbaum, 2005), and construct high-dimensional spaces using human-generated norms (e.g., Jamieson, Jain, Fernandez, Glattard, & Nowak, 2015; Kumar, Steyvers, & Balota, under review). Therefore, although models exist within this continuum of using language itself or responses from semantic tasks, a common difference between these models is the data used to define the representation.

The second distinctive issue pertains to *structure*—while feature-based and distributional models typically represent words in a *distributed* manner (e.g., as a list of features, along several dimensions, topics, etc.), networks typically represent words using a single node. Therefore, the types of processing assumptions that follow from the underlying structure have important implications for behaviors these models seek to explain. For instance, while networks offer a way to conceptualize similarity between concepts in terms of the "path length" between concept nodes, distributional and feature-based models consider the "angle" or the "overlap" between concept vectors to be indicative of similarity. Importantly, as mentioned, models can be transformed from being a high-dimensional space to a network (e.g., by thresholding cosine similarity values to construct "paths;" see Steyvers & Tenenbaum, 2005) and from a network to a high-dimensional space (e.g., by construction word association spaces; see Kumar et al., under review; Steyvers, Shiffrin, & Nelson, 2005). In addition, integrating feature-based information with distributional semantic representations has also been explored in the literature (e.g., Andrews, Vigliocco, & Vinson, 2009; Howell, Jankowicz, & Becker, 2005; Jones & Recchia, 2010). Given that different types of models can be flexibly modified to fit into one structure or another, the *data* used to build these models appear to be more critical than the underlying structure itself. The following section focuses on how networks generally created from behavioral norms have contributed to our understanding of semantic memory structure.

3. Semantic memory network structure

Although semantic networks have been critical in describing the structure of semantic memory, "network" is an umbrella term for several different ways of conceptualizing memory structure. As discussed, Collins and Quillian originally invoked a propositionally based hierarchical network, whereas Collins and Loftus' (1975) network was more associative in nature. Steyvers and Tenenbaum (2005) provided a starting point of using modern network science to understand the nature of semantic representation. They used three large datasets, the Nelson, McEvoy, and Schreiber (2004) free association norms, Roget's Thesaurus (Roget, 1911), and WordNet (Fellbaum, 1998; Miller, 1995) to construct three different semantic networks. Free association is a ubiquitous task in psycholinguistics where participants are asked to respond with one or several words that come to mind in response to a given cue word. Steyvers and Tenenbaum used responses produced from the Nelson et al.'s free association norms as a metric to connect edges between different word nodes (i.e., if at least two participants out of 120 participants on average responded *tiger* to *lion*, an edge was created between them; see

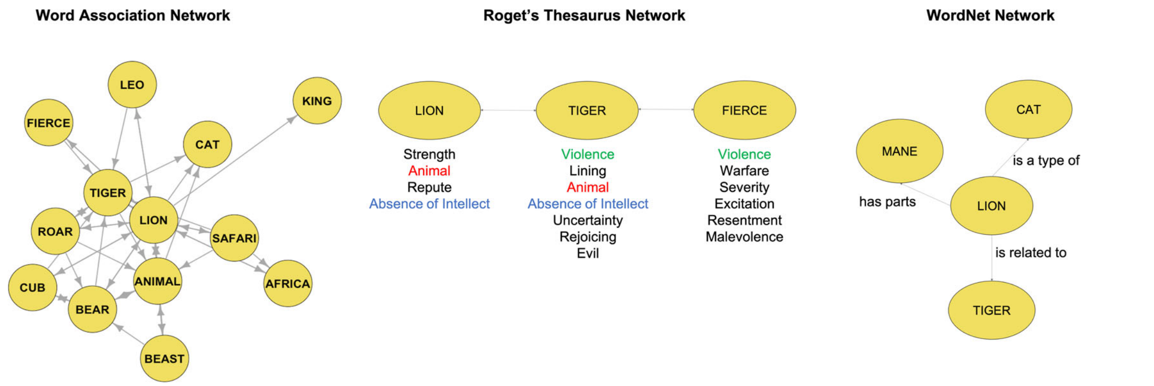

Fig. 2. Depiction of three semantic networks created by Steyvers and Tenenbaum (2005). LION and TIGER were connected in the word association network if TIGER was produced as a response to LION. In the thesaurus-based network, LION and TIGER were connected if they shared common semantic categories (e.g., *animal*). In the WordNet-based network, LION and TIGER were connected if there was a relationship between them (e.g., is related to).

left panel of Fig. 2). Roget's Thesaurus contains a large collection of words classified into different semantic categories, which Steyvers and Tenenbaum then used to connect edges between words by applying the criterion of at least one shared semantic category (second panel of Fig. 2). Finally, the WordNet-based semantic network captures connections between words and their meanings, and also different types of semantic relationships (e.g., antonym, hypernymy, etc.). Importantly, however, the *type* of relationship was not implemented within the Steyvers and Tenenbaum model, who instead used the presence or absence of any relationship between words as the criterion to draw edges within the network (see right panel of Fig. 2).

One important outcome of this network-based approach to representing meaning-related information is that it affords quantitative methods to estimate the *distance* between concepts as a function of number of intervening nodes (or path length), and ultimately test hypotheses for how semantic distance may guide behavior in cognitive tasks. For example, within an undirected associative network configuration (see Fig. 3), *lion* and *stripes* are 2 steps apart, compared to the directly connected *lion* and *tiger*, therefore providing the opportunity to test for mediated priming (see Balota & Lorch, 1986) and more distant semantic priming effects in priming-based tasks. Indeed, replicating Kenett, Levi, Anaki, and Faust (2017), Kumar, Balota, and Steyvers (2020) recently showed that path lengths from such networks not only accounted for the original mediated 2-step priming effects, but also differences between more distant concepts, that is, there was a reliable difference in response latencies in a primed perceptual identification task for words representing concepts that were four versus six steps away within associative semantic networks.

Although the number of links between two concepts provides a measure of distance, it is important to acknowledge that the notion of path length is somewhat arbitrary. Specifically, network construction requires an explicit decision on how to define an edge between two concepts (see Castro & Siew, 2020, for a detailed discussion). As noted, in the Steyvers and

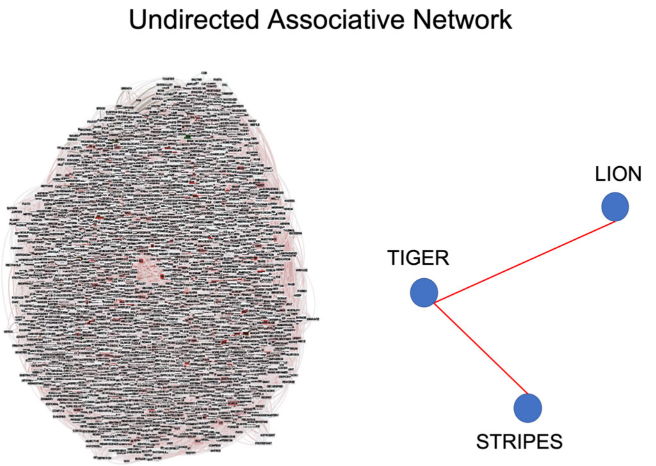

Fig. 3. Large-scale structure and path from LION to STRIPES in the undirected associative network created by Steyvers and Tenenbaum (2005), adapted from Kumar et al. (2020).

Tenenbaum work, they implemented an edge in their associative network if a cue (e.g., *lion*) produced the same first response (e.g., *tiger*) for at least two individuals out of 120 individuals on average. Why not 1, 10, or 20 responses out of 120 to produce an edge? We will refer to this as the *all-or-none* problem. Another approach would be to directly implement strength of association (as defined via free association norms) as the "distance" between any two nodes, as in the original Collins and Loftus (1975) model. There are clear effects of strength of association and category dominance in priming tasks (e.g., Balota & Duchek, 1988; Lorch, 1982) that would support this type of structural organization. Indeed, more recent work has relied on the frequency of response production as a measure of edge weight in association networks (e.g., De Deyne, Perfors, & Navarro, 2016; De Deyne & Storms, 2008), and has been shown to account for similarity and relatedness judgments. Other researchers have used a technique called *percolation analysis* to examine whether systematically removing edges within a network with weights/strengths below a certain threshold can provide useful information about the resilience, flexibility, and robustness of individual semantic networks (Cosgrove, Kenett, Beaty, & Diaz, 2021; Kenett et al., 2018; also see Gruenenfelder, Recchia, Rubin, & Jones, 2016. for a DSM-based network thresholding study on free association). Therefore, although this flexibility in edge construction within semantic networks can be viewed as problematic from an interpretation standpoint, systematically measuring how human behavior maps on to different network-based thresholds could also provide novel insights into individual differences in concept connectivity and cognitive flexibility.

3.1. Structural properties and growth in semantic networks

Another important aspect of modern semantic network analyses is the opportunity to examine small- and large-scale properties of the network using graph-theoretical measures. These include node degree (referring to the total number of connections for a given node in a network) and clustering coefficient (a metric that describes the extent to which the neighbors of a node are interconnected and is an indicator of network density), among others. These metrics have been useful in characterizing the semantic networks of typical and atypical populations (e.g., see Christensen, Kenett, Aste, Silvia, & Kwapil, 2018; McNally et al., 2015), as well as examining patterns related to creativity (e.g., Kenett & Faust, 2019; Kenett, Anaki, & Faust, 2014) and aging (Wulff, De Deyne, Jones, Mata, & Aging Lexicon Consortium, 2018).

Steyvers and Tenenbaum found that the semantic networks they implemented exhibited *small-world structure*, which is also a property of naturally occurring networks such as the World Wide Web, as reflected by sparse connectivity, short average path lengths, and strong clustering. This leads to the important question of *why* naturally occurring networks would develop these characteristic signatures. Barabási and Albert (1999) proposed that new connections in a network are made in proportion to the existing number of connections for a given node, therefore following a "rich get richer" principle, formally called the principle of *preferential attachment*. Finding support for this hypothesis, Steyvers and Tenenbaum (2005) modeled networks based on existing connections and showed that these networks closely resembled behavioral associative networks. Other work in this domain has examined whether the growth of semantic networks among children actually follows these patterns. For example, using a longitudinal sample of nouns acquired by children of ages 16–30 months, Hills, Maouene, Maouene, Sheya, and Smith (2009) showed that there was stronger evidence for *preferential acquisition* among children, instead of preferential attachment. According to preferential acquisition, words are acquired in proportion to how well they connect to other words in the *learning environment*, not how well they connect to already-learned words in the internal lexicon, as would be predicted by the preferential attachment hypothesis (also see Fourtassi, Bian, & Frank, 2020). Overall, these growth mechanisms are critical in understanding how new words may be acquired from the natural environment by children as well as how cognitive structures and the processes involved in language acquisition are inextricably linked to each other.

3.2. Relational information within semantic networks

Another important issue in current network science is the extent to which specific networks can capture the varied contexts within which a word may be found or used. Consider the word *lion*. Within a directed associative network (Steyvers & Tenenbaum, 2005), *lion* has 12 direct outgoing links, as shown in Fig. 1 (first panel), which represent different types of relationships (e.g., a *lion* <can be spotted at a> *safari* vs. a *lion* <is similar to a> *tiger*) but are not classified as such in a nonlabeled network. This example demonstrates the *diversity* as well as *levels* of meaning-related information available to humans within their semantic memory, which may not be fully captured in commonly implemented semantic networks. Therefore, there is a need to supplement semantic networks created from semantic tasks with explicit relational information to accurately reflect the structure and interconnectedness of semantic memory.

There have recently been some promising attempts to build networks that include specific types of relational links between words. ConceptNet (Speer, Chin, & Havasi, 2016) is a prime example of such a development. ConceptNet links concepts through different types of relationships, as opposed to word association networks where there is no distinction between the types of edges. This network configuration is based on a number of crowdsource databases including Open Mind Common Sense (Singh, 2002) and "Games with a purpose" designed to collect common knowledge (Nakahara & Yamada, 2011). Edges in ConceptNet are labeled according to 36 different types of relations. For example, symmetric relations might include "located near" and "similar to," whereas asymmetric relations would include "capable of" and "created by." Importantly, Speer et al. not only constructed this edge-labeled network but also used a retrofitting procedure to combine ConceptNet with *distributional* word representations from two popular machine learning-based DSMs, word2vec (Mikolov et al., 2013) and GloVe (Pennington et al., 2014), to create ConceptNet Numberbatch, a more powerful semantic representation model. ConceptNet Numberbatch outperformed state-of-the-art DSMs on a series of cognitive tasks including semantic relatedness judgments, analogy completion, and story completion. This is a significant step forward in the field, which suggests that one may need to include labelled pathways connecting concepts to represent semantic memory structure. Hybrid models utilizing distributional semantics from large language corpora may be particularly advantageous to this enterprise.

ConceptNet Numberbatch represents an important shift in reconciling network-based accounts with alternative representational accounts of semantic memory structure. However, it also raises important questions about how search and retrieval would operate within such a configuration. Would one search via the specific types of relationships or the specific concepts presented in a task? Alternatively, there may be an initial fast nonspecific search process reflecting global similarity followed by a more detailed search process. This highlights the importance of considering the types of *processes* demanded by a task to access specific relations within a network.

4. Retrieval processes from semantic memory

As discussed earlier, structure always needs to be coupled with processes that act on that structure to perform a specific task, given that different tasks differentially emphasize different components of lexical/semantic processing (see Balota, Paul, & Spieler, 1999, for a discussion). Recent work in semantic memory research has focused on computationally implementing processing frameworks that could account for search and retrieval processes within semantic memory. The foundational ideas of spreading activation proposed by Collins and Loftus (1975, also see Anderson, 1983) discussed earlier, have helped guide some of the research in this domain. Specifically, the notion of a word activating its neighbors in proportion to the associative strength between the words has been formalized via *stochastic random walks* within a semantic network. Random walk models assume that a "walker" starts from a particular node and selects subsequent nodes to traverse within a network based on the edge weight (or similarity) between the starting node and its neighbors (see Fig. 4).

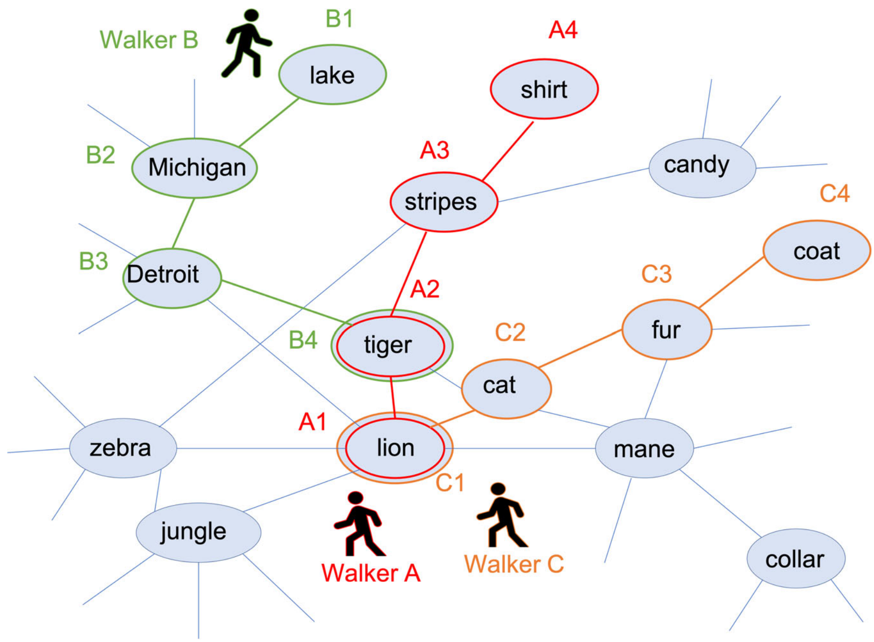

Fig. 4. A schematic of random walks within a semantic network. Walkers A and C both start from LION and traverse different paths. Walker B starts from LAKE and ends the walk on TIGER.

Within these random-walk models, edge weights are typically derived from behavioral responses (e.g., free association norms) and the parameters that govern the walk can provide powerful insights into how an individual may search memory. For example, a decay parameter that penalizes longer paths over shorter paths has been shown to be important in certain tasks such as relatedness judgments (De Deyne, Verheyen, et al., 2016). Overall, principles of random walks have been successfully applied to explain performance in several semantic tasks, such as semantic fluency (Zemla & Austerweil, 2018), letter fluency (Griffiths, Steyvers, & Firl, 2007), and the Remote Associates Test (Bourgin, Abbott, Griffiths, Smith, & Vul, 2014; Smith, Huber, & Vul, 2013), as well as to explore the memory structure of creative individuals (Kenett & Austerweil, 2016).

Another domain where random walks and path-based models have been applied to understand complex search and retrieval processes is within the context of language games (e.g., Beckage, Steyvers, & Butts, 2012; Fathan et al., 2018; Kumar et al., under review). For example, Kumar et al. (under review) applied random-walk-based models to examine search and retrieval within a cooperative word game, Connector. Within Connector, a Speaker is given two words (*lion* and *tiger*) from a 20-item word board and asked to produce a one-word clue (e.g., *cat*) that is related to both words (see Fig. 5). The Guesser then attempts to identify the two words on the board that the Speaker might have been referring to, based on the clue (e.g., given the clue *cat*, a Guesser might select *lion-snake* from the board). Kumar et al. showed how player responses were predicted by a distance measure capturing random

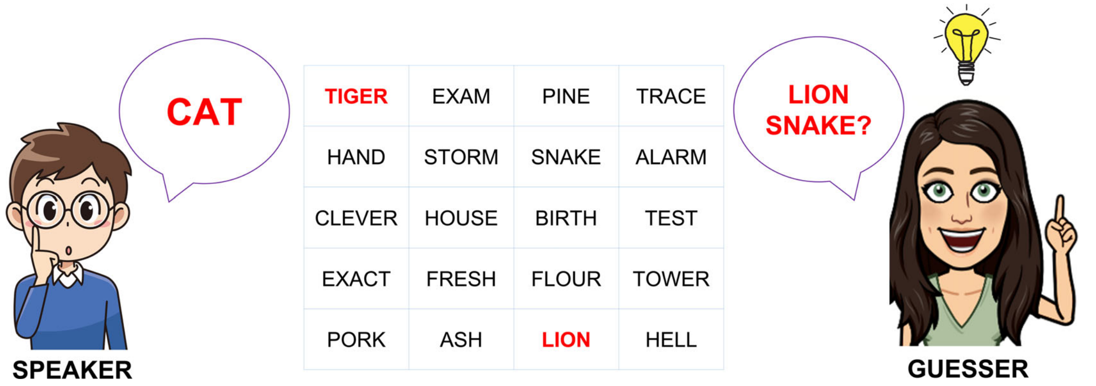

Fig. 5. An experiment trial in Connector in Kumar et al. (under review). The Speaker generates a one-word clue (CAT) to a word-pair on the board (LION-TIGER). The Guesser responds with two words (LION-SNAKE).

walks operating over an associative semantic network. Of course, this is not an explicit search model and future work should examine how local and global search processes and/or foraging can account for unconstrained search processes within semantic memory.

In addition to implementing random walks, some recent proposals have attempted to quantify how activation from a given node may eventually spread within the network across several time steps. For example, Vitevitch, Ercal, and Adagarla (2011) and Siew (2019; also see De Deyne, Verheyen, et al., 2016, for a similar model) formalized the notion of spreading activation in a model based on connections within either phonological or associative networks, respectively, and showed that these activations successfully explained performance in spoken word recognition, false memory recall, as well as large databases of semantic priming effects in lexical decision and naming tasks (Hutchison et al., 2013). Of course, it is again important to consider how this conceptualization of spreading activation would fit with a relationally organized lexicon (e.g., ConceptNet), that is, would activation spread to relation-specific nodes, or diffuse through the complete network? If activation does, in fact, differentially spread across different relationships, are some relationships more predictive of performance in priming and fluency tasks than others? Although there are several remaining questions to be answered, this work is clearly important in illustrating how a network-based structural approach can be effectively combined with a temporally constrained process (spreading activation) to explain behavior.

Random walks represent one proposal for retrieval from semantic memory, and alternative processing mechanisms have also been explored in the literature. For example, Hills, Jones, and Todd (2012; also see Hills et al., 2015) proposed a two-process model operating over a distributional model of semantic memory (BEAGLE; Jones & Mewhort, 2007) to account for behavior in the semantic fluency task. Their model assumes a *local* search process focused on random-walk-type processes operating over the semantic space to produce semantically related clusters of items, and a *global* process focused on long-distance “jumps” within semantic space when the local cluster is sufficiently depleted, similar to animal foraging patterns found in the wild (Hills, Kalff, & Wiener, 2013). This two-stage local-global switching process has surface-level similarity to the two-process decision model proposed by Smith et al. (1974), wherein a fast familiarity-based process is followed by a slower analytic process. Indeed, it is possible that more global semantic information is retrieved relatively quickly, which is followed by a more detailed search of specific relations, possibly driven by labeled edges. The distinction between fast automatic nonanalytic processes followed by more attention-driven analytic processes is fundamental to a wide variety of two-stage models of human behavior (see, Atkinson & Juola, 1974; Balota & Chumbley, 1984; Jacoby, Jones, & Dolan, 1988). Investigating how this two-process distinction fits within a network-based perspective to fully account for human behavior in semantic tasks is an important next step for semantic network research.

5. Semantic networks: Looking under the hood

Although we have primarily focused on a network-based approach, as noted there are important alternative models of semantic memory (e.g., DSMs), and indeed, there is empirical work comparing the extent to which different classes of models capture human behavior. For example, Vankrunkelsven, Verheyen, Storms, and De Deyne (2018) compared a DSM trained on text corpora to an associative model based on the Small World of Words (De Deyne et al., 2019) free association dataset, and showed that the associative model significantly outperformed the DSM in predicting behavioral ratings of words on lexical properties such as age of acquisition and valence. Other work has also shown that associative models generally better capture relatedness and similarity judgments (De Deyne, Perfors, et al., 2016; De Deyne, Verheyen, et al., 2016), conceptual association in language games (Kumar et al., under review), and visual and affective features of concepts (De Deyne, Navarro, Collell, & Perfors, 2021), compared to DSMs.

There are two possible reasons why associative network models outperform DSMs. One possibility is that associative models are more likely to capture perceptual/motor aspects of meaning that are missing from DSMs created from textual corpora. However, it is a priori unclear why associative norms might capture such sensory/motor information any more than natural language. This issue is related to the "grounding" problem in cognitive science, wherein DSMs have been previously criticized for solely relying on text corpora to develop semantic representations (Barsalou, 2016), in addition to being strongly biased by factors such as the size of the text corpora and parameter tuning (Kenett, 2019). Indeed, accumulating evidence suggests that semantic memory is more likely a combination of linguistic, affective, and sensorimotor interactions (for a review, see Kumar, 2021b). To reflect this emerging view, multimodal DSMs (e.g., Bruni et al., 2014) and multiplex networks (e.g., Stella, Beckage, & Brede, 2017) that integrate different sources of information (perceptual, phonological, etc.) to construct semantic representations are an active area of research in this field.

However, another possibility is that the success of semantic networks at accounting for behavior in cognitive tasks is entirely due to the *shared variance* between the free association task and the tasks on which DSMs and associative models are compared. Specifically, free association and fluency data are among the most popular methods to create semantic networks (e.g., Christensen and Kenett, 2019; De Deyne et al., 2008; Kenett et al., 2017; Nelson, Dyrdal, & Goodmon, 2005; Zemla & Austerweil, 2018). It is important to remember that free association norms and semantic fluency data represent outcomes of semantic retrieval *tasks* in and of themselves. Therefore, one may be concerned about a type of circularity, that is, there may be considerable overlap in the retrieval processes used in these norming tasks (e.g., free association) and other semantic tasks, which may produce an advantage for behavioral networks in accounting for performance in these tasks (see Jones, Hills, & Todd, 2015, for detailed arguments). Indeed, comparisons of DSMs versus associative models that differ *solely* on the underlying data source and use similar metrics (e.g., cosine similarity) for both types of models (e.g., De Deyne, Verheyen, et al., 2016; Kumar et al., under review) continue to demonstrate advantages for associative models. Furthermore, even when DSMs and associative models are *both* supplemented with additional visual and/or affective information as in De Deyne et al. (2021), associative models still continue to better capture behavioral performance across relatedness/similarity judgments and tasks that rely on such information. Therefore, the *data* underlying network models (free association) and distributional models (natural language corpora) appear to be critical when considering their relative predictive power, and the shared variance between free association and other tasks may be confounding some of these observed patterns.

An important insight from these comparative studies is that there are strengths and limitations to utilizing human-generated norms versus natural language corpora to create semantic representation models. First, it is at least possible that free associations may evoke certain types of meaning-related processing (e.g., mental imagery, emotional experiences, etc.) that is not consistently reflected in natural language corpora. For example, the mental image of a *banana* being *yellow* or being *peeled* is extremely salient and *yellow* and *peel* may therefore be frequently produced as responses to *banana* within a free association task,[3] whereas DSMs that capitalize on co-occurrence patterns in text may instead emphasize the similarity of a *banana* to other fruit such as *apple* and *mango*.[4] In this light, network models based on free associations may therefore provide *complementary* information to text-based DSMs. Second, it is certainly possible that there is some shared variance between the free association task and the evaluation tasks, although this could not easily explain why DSMs sometimes outperform associative models in semantic priming tasks (e.g., Kumar et al., 2020), or why the benefit of free association data over DSMs is minimal for some types of lexical information (e.g., concreteness, Vankrunkelsven et al., 2018). Ultimately, it is likely that the success of free association-based models is a combination of *some* shared variance, and also meaningful semantic information not presently captured within DSMs. Hence, future work should focus on (a) better understanding the processes underlying free association (as in Gruenenfelder et al., 2016; Kumar, 2021a; Nelson, McEvoy, & Dennis, 2000; Richie, Aka, Bhatia, in preparation), (b) using associative network models as an *evaluative baseline* for DSMs (see Kumar et al., under review), and (c) exploring "network"-based models with different underlying sources of information (e.g., multiplex networks, text-based networks, etc.).

Ultimately, the strength of the semantic "network" approach may be considered independent of the underlying data source. Regardless of whether the data come from human-generated norms (as has been typical), distributional models (as in Gruenenfelder et al., 2016), thesauri, or WordNet (as in Steyvers & Tenenbaum, 2005), one can still investigate questions about *structure* and *process* using the network approach. Critically, it is important to conduct studies that systematically manipulate the underlying data source within different network-based models to understand which type of data best captures human performance across different semantic tasks. Indeed, as shown by Steyvers and Tenenbaum (2005), networks created from different data sources can show similar large-scale *global* connectivity, but still differ in critical *local* hubs (highly connected nodes), as well as the extent to which they map onto behavioral data from naming and lexical decision tasks. Additionally, Gruenenfelder et al. (2016) showed that a *hybrid* DSM model that incorporated *both* associative (direct co-occurrence-based) and contextual (indirect co-occurrence-based) information from text-based corpora best captured different properties of free association norms (e.g., power distributions, clustering, etc.). Therefore, meaning may be represented at different *levels* (contextual and associative), and the network approach allows researchers to effectively test different models and identify different aspects of what constitutes meaning at the *global* and *local* scale.

6. Current questions in semantic network research

Given the promise of network-based approaches to examine semantic memory structure and processes, it is also important to consider outstanding questions in the field, and the extent to which networks are capable of answering these questions. This section will focus on the strengths and potential limitations of the network-based approach, and also provide future directions aimed at improving our overall understanding of semantic structure and processing.

6.1. Automatic spreading activation versus attentional retrieval processes

As noted, whenever one considers process-based assumptions regarding network models, it is important to consider the extent to which activation and search processes are under the control of attention. In a landmark semantic priming study, Neely (1977) was able to totally dissociate an automatic fast-acting spread of activation from a slower attention-demanding retrieval process. For example, the word *cat* could be automatically activated by the brief presentation of *dog*, via an automatic spreading activation mechanism, but one could also obtain priming for *cat* if participants engage in a controlled process of predicting which words might follow *dog*. This original dissociation nurtured considerable empirical work that fit within this perspective (see McNamara, 2005; Neely, 1991, for a review of priming studies), and the distinction between automatic processes and attention-demanding processes is central to virtually all aspects of cognition from pattern recognition (see Treisman, 1969), memory retrieval (see Jacoby et al., 1988), to person perception (see Smith & DeCoster, 2000). As noted, it is possible that qualitatively distinct information (e.g., associative vs. relational) or different characteristics of the same network (e.g., local vs. global) may be retrieved via these different retrieval modes. Ultimately, we believe that any model of structure and process will need to incorporate such distinct retrieval processes to fully account for observed behavior in a given task, or at the very least demonstrate how a single process model may accommodate the data that have been used to support multiple process models.

6.2. Compositionality

An important aspect of language is that it is compositional, that is, phonemes combine to form words, words combine to form phrases, and so on. Although semantic networks have been mostly conceptualized and applied at the lexical/conceptual level, some recent work has examined how words or concepts can be combined to form higher level conceptual structures. For example, Kenett and Thompson-Schill (2020) analyzed how semantic network parameters changed in a conceptual combinations task, where participants were given ambiguous noun pairs (e.g., *robin-hawk*), and produced interpretations that were *attributive* (applying a dominant feature of one word to another, for example, a *red-breasted* robin-like *hawk*) or *relational* (connecting two words thematically, e.g., a *hawk* that preys on *robins*). The authors found that the associative networks constructed after the conceptual task exhibited greater global connectivity (i.e., shorter average path lengths and higher global clustering) when participants were primed with *relational* interpretations. This work highlights the dynamic nature of semantic knowledge, and more importantly, provides insights into how concepts represented at the word level may be combined to produce higher order compositional structures. Of course, more work is needed to understand how these higher order structures that emerge during tasks can be computationally realized within a network model of semantic memory. For example, how might a sentence (e.g., *Mary loves Jack*) be represented using a semantic network, and how might syntax influence this representation (e.g., distinguishing between *Mary loves Jack* and *Jack loves Mary*)?

At present, network models lack an account of how representations scale up to truly represent these higher order structures, whereas some DSMs propose that individuals use prediction error and linguistic context to guide their internal representations for sentences, phrases, and even events (e.g., Elman, 1990; Franklin, Norman, Ranganath, Zacks, & Gershman, 2020; for a review, see Kumar, 2021b). In addition, there are studies showing that some type of vector composition or compounding process applied over distributional semantic representations results in compositional representations (e.g., Marelli, Gagné, & Spalding, 2017; Mitchell & Lapata, 2010). Therefore, it is important to reconcile the distributional approach with the network-based approach. The elaborate labeled connections within ConceptNet Numberbatch discussed earlier are an important step in this direction. Of course, this nicely dovetails with the previous section, because it is quite possible that the higher order aspects of comprehension may demand qualitatively distinct retrieval processes that may be more attention-demanding, compared to more automatic activation processes.

6.3. Newer insights from machine learning

In addition to exploring how DSMs trained on language corpora explain human behavior, there has also been a growing interest in machine learning to extract *relational* structures between concepts. Specifically, although DSMs typically represent words in a high-dimensional space and provide a natural way of estimating semantic "distance" in this space (e.g., how related are *lion* and *zebra*, compared to *lion* and *safari*), the specific *types* of relationships encoded within these representations remain unknown (e.g., *lion* and *zebra* can both be spotted at a safari, whereas only a *lion* is a *predator*). Some recent work has attempted to address this question by incorporating relational learning during word vector computations. For example, Camacho-Collados, Espinosa-Anke, and Schockaert (2019) used a neural network model to *learn* distinct relational word vectors for different words within a sentence and showed that for a given word pair (e.g., *innocent-naïve*), their relational representations identified subtle relationships as indicated by nearest neighbors (e.g., *vain-selfish*, *cruel-selfish*), compared to the nearest neighbors produced by a nonrelational DSM that were harder to interpret (e.g., *murder-young*, *conspiracy-minded*; also see Jameel et al., 2018). Constructing *relational* representations echoes ideas similar to encoding relational information via ConceptNet discussed earlier, but goes a step further in that it does not rely on crowdsourced knowledge bases but explicitly derives these relations from text. This method has the potential to uncover abstract and subtle relationships that may not be encoded well enough through specific relation categories. This is clearly an important development that could potentially be applied to construct labeled semantic networks, which, in turn, could have broad implications for how well relational networks account for human performance. A promising future direction would be to examine how well semantic networks constructed from relational vector representations encode aspects of human behavior observed in cognitive tasks, a domain that remains relatively understudied.

Network or graph-based approaches to knowledge representation have also been explored from a machine learning perspective (for a review, see Grohe, 2020). Specifically, in recent years, there has been considerable interest in identifying latent communities within real-world networks using machine learning techniques. For example, Perozzi, Al-Rfou, and Skiena (2014) proposed DeepWalk, an algorithm that combined the notion of a random walk with inferring the next words in a sequence using insights from the predictive word2vec model. Perozzi et al. showed how DeepWalk successfully captured social relationships in a network of bloggers and Youtube users (also see node2vec; Grover & Leskovec, 2016). Other research in this area has focused on training more sophisticated neural networks, such as graph neural networks (GraphSage; Hamilton, Ying, & Leskovec, 2017) that are able to capture more significant changes to the graph structure, such as the addition of new nodes, and deep network embeddings (Line; Tang et al., 2015). These networks preserve the local and global structure to produce meaningful representations that can then be applied to downstream tasks such as predicting upcoming links in a sequence of nodes. The task of link prediction is particularly interesting, because it allows researchers to probe how information may propagate within a network, and echoes ideas similar to spreading activation discussed earlier. This work is important in illustrating how existing *semantic* network-based accounts can be further strengthened by borrowing techniques from machine learning. Specifically, uncovering latent dimensions from associative semantic networks using algorithms like DeepWalk or GraphSage could provide further insights into how concepts cluster into meaningful neighborhoods as well as change as a function of new information entering the network. Furthermore, utilizing algorithms of link prediction could provide mechanisms to develop compositional representations at the phrase or sentence level within semantic networks, which could ultimately inform our understanding of how semantic knowledge accumulates over time and guides behavior.

6.4. Semantic networks within the brain

The discussion so far has conceptualized semantic network models as a viable account for how concepts are organized and retrieved to perform different tasks. However, a basic assumption in all of this work is that we are somehow approximating the mind, which ultimately needs to be implemented in a physical brain. Although early explorations into connectionist modeling and neural networks (e.g., Rumelhart, Hinton, & McClelland, 1986) were inspired by neurobiology, state-of-the-art neural network models are not considered accurate models of the brain (see Bengio, Goodfellow, & Courville, 2017, for a discussion). In a separate domain of research, intrinsic brain networks (e.g., Fox, Zhang, Snyder, & Raichle, 2009) are considered central to distinct types of cognitive processing. It is important to emphasize here that although "networks" is used as a catch-all term across different fields, it represents qualitatively and quantitatively different information from the types of *semantic* networks discussed in this review. For example, *neural* networks delineate error-driven learning mechanisms for capturing information across different sets of nodes, whereas *brain* networks are based on regions of the brain that are activated/inhibited during task performance. Both types of networks are in contrast to *semantic* networks that mainly consider words as the unit of interest. This begs the question of whether *semantic* networks are neurobiologically plausible accounts of meaning representation. One way to answer this question may be to conceive meaning as a multivoxel neural pattern of activity (e.g., Fedorenko, Nieto-Castanon, & Kanwisher, 2012; also see Musz & Thompson-Schill, 2017), and semantic networks could then potentially track the flow and structure of information activated across these time-dependent distinct neural patterns. Although there is some evidence that semantic retrieval may indeed be represented via distinct neural activity (e.g., Musz & Thompson-Schill, 2015), such neurally inspired semantic network accounts remain relatively underexplored. Another possible way to reconcile semantic networks with neurobiological accounts may be to assume that semantic networks exist at the *algorithmic* or *computational* level, whereas neurobiological accounts of semantic memory exist at the *implementational* level, consistent with Marr's (1982) levels of analysis. In this way, a complete theory of semantic memory may emerge from eventually integrating the algorithmic and conceptual accounts of meaning-based processing with a structural implementation at the neuronal level.

7. Discussion

The notion of networks has come very far within the span of over half a century when it was first proposed as a model for human semantic memory (Collins & Quillian, 1969; Quillian, 1967). As this review has described, the field has progressed from purely conceptual accounts of meaning representation to quantitative methods of estimating the size and nature of semantic memory. At the same time, the progress in computational modeling has led to deeper explorations into the ways in which humans learn and access such networks, giving rise to more complex models of semantic knowledge and processing, as well as raising questions about the validity and cognitive plausibility of these models. The present review focused on how cognitive network science has been instrumental in describing the *structure* of semantic knowledge as well as the cognitive *processes* that guide navigating this structure to ultimately produce behavior. These discussions have emphasized the utility of the network approach to capture the flexibility of semantic relationships as well as temporal processes that govern retrieval from such networks.

As we have seen, the term "networks" can refer to a broad range of models, ranging from associative semantic networks and knowledge graphs to social or biological networks. These networks differ both in the data contributing to the *structure* of the network, as well as the underlying *processes* operating upon this structure. In the present review, we have described how networks created from behavioral norms have been useful in explaining a wide array of cognitive phenomena. However, we have also noted some potential issues regarding circularity within this approach, and emphasized that purely associative networks may not be sufficient to fully characterize the vast amount of relational information available to humans. We have also provided future directions to integrate this information into existing networks by utilizing distributional models (e.g., ConceptNet Numberbatch) and techniques being widely applied in machine learning (e.g., Camacho-Collados et al., 2019). Similarly, in thinking about the *processes* that drive the network-based perspective, as discussed, the notion of random walks, automatic spreading activation versus attentional analytic processes, and growth mechanisms have all been applied to explain a wide range of cognitive phenomena. In fact, as the preceding section highlighted, machine learning researchers are now leveraging tools from graph theory to adequately represent and study real-life networks. Therefore, an integrated account of semantic memory may indeed involve extracting distributional information from natural environment and ultimately representing this information using network-based structures that may provide insights to more general dynamics of cognitive processing.

In light of the work discussed, one way to reconceive semantic networks may be to use natural language (the proxy for which is typically large text corpora) and other nonlinguistic sources of information (such as images, phonology, and affective stimuli) as a starting point from which relational and concept learning occurs, and then augment this learning process with processing mechanisms that directly follow from a network-based perspective. For example, one could conceptualize a *multimodal* DSM that also learns relational information (via text corpora or knowledge graphs) to produce semantic representations that can ultimately be used to guide processing via spreading activation or random walks. Thus, an important takeaway from this review is to potentially envision semantic networks as a collection of *tools* and *methods* to conduct deeper explorations into the structure of knowledge representation.

Ultimately, the goal of semantic memory research is to develop a computational model of knowledge structure and the processes that describe the different ways in which individuals learn and use meaning. So, an obvious question may be whether semantic networks are a *literal* account of human knowledge and processing, that is, does human knowledge truly exist as a network? Although explorations into exactly how knowledge is represented within the brain will continue (see Castro & Siew, 2020, for a discussion), in light of the work discussed in this review, it seems most likely that modern distributional models (specifically multimodal DSMs) provide a promising account of *learning* meaning from natural environment, whereas semantic network accounts provide useful *conceptual* tools to probe these representations and the processes that operate upon these representations. A simple analogy from mathematics may be helpful here: although it is impossible to show the existence of a negative number, we can all agree that the *idea* of a negative number indicating the lack of something is clearly useful (e.g., -1 degree means colder weather) and has widespread applications and implications in fields such as psychology, geography, and physics. In a similar manner, although the *physical* existence of semantic networks demands further inquiry, the present review has provided substantial evidence in favor of considering the *conceptual* and *computational* utility of envisioning semantic memory as a network, which ultimately allows us to ask and answer more focused questions about how human knowledge guides cognition.

Notes

1 Computational feature-based models have significantly contributed toward understanding semantic memory organization and processing (see Cree & McRae, 2003; McRae, 2004; McRae, De Sa, & Seidenberg, 1997), and have also been recently integrated with quantitative semantic network-based accounts (Solomon, Medaglia, & Thompson-Schill, 2019) to account for flexible concept use. However, despite some advantages of feature-based models (see Buchanan, Valentine, & Maxwell, 2019), there is also some concern regarding how the critical features for a concept are learned in the first place (see Jones, Willits, Dennis, & Jones, 2015). Due to space limitations, we will not fully explore feature-based representations in the current paper.

2 It is important to note here that there are also other types of semantic models, such as dynamic attractor networks (e.g., McLeod, Shallice, & Plaut, 2000), although we explicitly focus on DSMs and semantic network models in the current paper.

3 The empirical frequencies of producing *yellow*, *peel*, and *apple* in response to *banana* in the Small World of Words database are 65, 16, and 10, respectively, see https://smallworldofwords.org/en/project/explore

4 See http://bionlp-www.utu.fi/wv_demo/ for a fast demo on nearest neighbors of words within a popular text-based DSM (word2vec).

References

Anderson, J. R. (1983). A spreading activation theory of memory. *Journal of Verbal Learning and Verbal Behavior*, *22*(3), 261–295.

Andrews, M., Vigliocco, G., & Vinson, D. (2009). Integrating experiential and distributional data to learn semantic representations. *Psychological Review*, *116*(3), 463.

Atkinson, R. C., & Juola, J. F. (1974). Search and decision processes in recognition memory. In D.H. Krantz, R.C. Atkinson, R.D. Luce, & P. Suppes (Eds.), *Contemporary developments in mathematical psychology: I. Learning, memory and thinking*. San Francisco: WH Freeman.

Balota, D. A., & Chumbley, J. I. (1984). Are lexical decisions a good measure of lexical access? The role of word frequency in the neglected decision stage. *Journal of Experimental Psychology: Human Perception and Performance*, *10*(3), 340–357.

Balota, D. A., & Duchek, J. M. (1988). Age-related differences in lexical access, spreading activation, and simple pronunciation. *Psychology and Aging*, *3*(1), 84–93.

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), 336–345.

Balota, D. A., Paul, S. T., & Spieler, D. H. (1999). Attentional control of lexical processing pathways during word recognition and reading. In S. Garrod & M. Pickering (Eds.), *Language Processing* (pp. 15–57). London: Psychology Press.

Barabási, A. L., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, *286*(5439), 509–512.

Barsalou, L. W. (2016). On staying grounded and avoiding quixotic dead ends. *Psychonomic Bulletin & Review*, *23*(4), 1122–1142.

Beckage, N., Smith, L., & Hills, T. (2011). Small worlds and semantic network growth in typical and late talkers. *PLoS One*, *6*(5), e19348.

Beckage, N., Steyvers, M., & Butts, C. (2012). Route choice in individuals—semantic network navigation. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 34, No. 34).

Bengio, Y., Goodfellow, I., & Courville, A. (2017). *Deep learning* (Vol. *1*). Cambridge, MA: MIT Press.

Bourgin, D., Abbott, J., Griffiths, T., Smith, K., & Vul, E. (2014). Empirical evidence for Markov chain Monte Carlo in memory search. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. *36*, No. 36).

Bruni, E., Tran, N. K., & Baroni, M. (2014). Multimodal distributional semantics. *Journal of Artificial Intelligence Research*, *49*, 1–47.

Buchanan, E. M., Valentine, K. D., & Maxwell, N. P. (2019). English semantic feature production norms: An extended database of 4436 concepts. *Behavior Research Methods*, *51*(4), 1849–1863.

Camacho-Collados, J., Espinosa-Anke, L., & Schockaert, S. (2019). *Relational word embeddings*. arXiv preprint arXiv:1906.01373.

Castro, N., & Siew, C. S. (2020). Contributions of modern network science to the cognitive sciences: Revisiting research spirals of representation and process. *Proceedings of the Royal Society A*, *476*(2238), 20190825.

Christensen, A. P., & Kenett, Y. N. (2019). *Semantic network analysis (SemNA): A tutorial on preprocessing, estimating, and analyzing semantic networks*. PsyArXiv.

Christensen, A. P., Kenett, Y. N., Aste, T., Silvia, P. J., & Kwapil, T. R. (2018). Network structure of the Wisconsin schizotypy scales–short forms: Examining psychometric network filtering approaches. *Behavior Research Methods*, *50*(6), 2531–2550.

Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, *82*(6), 407–428.

Collins, A. M., & Quillian, M. R. (1969). Retrieval time from semantic memory. *Journal of Verbal Learning and Verbal Behavior*, *8*(2), 240–247.

Cosgrove, A. L., Kenett, Y. N., Beaty, R. E., & Diaz, M. T. (2021). Quantifying flexibility in thought: The resiliency of semantic networks differs across the lifespan. *Cognition*, *211*, 104631.

Cree, G. S., & McRae, K. (2003). Analyzing the factors underlying the structure and computation of the meaning of chipmunk, cherry, chisel, cheese, and cello (and many other such concrete nouns). *Journal of Experimental Psychology: General*, *132*(2), 163–201.

De Deyne, S., Navarro, D. J., Perfors, A., Brysbaert, M., & Storms, G. (2019). The “Small World of Words” English word association norms for over 12,000 cue words. *Behavior Research Methods*, *51*(3), 987–1006.

De Deyne, S., Navarro, D. J., Collell, G., & Perfors, A. (2021). Visual and affective multimodal models of word meaning in language and mind. *Cognitive Science*, *45*(1), e12922.

De Deyne, S., Perfors, A., & Navarro, D. J. (2016, December). Predicting human similarity judgments with distributional models: The value of word associations. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers* (pp. 1861–1870).

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

De Deyne, S., & Storms, G. (2008). Word associations: Network and semantic properties. *Behavior Research Methods*, *40*(1), 213–231.

De Deyne, S., Verheyen, S., & Storms, G. (2016). Structure and organization of the mental lexicon: A network approach derived from syntactic dependency relations and word associations. In A. Mehler, P. Blanchard, B. Job, & S. Banish (Eds.), *Towards a theoretical framework for analyzing complex linguistic networks* (pp. 47–79). Berlin: Springer.

Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, *14*(2), 179–211.

Fathan, M. I., Renfro, E. K., Austerweil, J. L., & Beckage, N. M. (2018). Do humans navigate via random walks? Modeling navigation in a semantic word game. In *Proceedings of the Annual Meeting of the Cognitive Science Society*.

Fedorenko, E., Nieto-Castanon, A., & Kanwisher, N. (2012). Lexical and syntactic representations in the brain: An fMRI investigation with multi-voxel pattern analyses. *Neuropsychologia*, *50*(4), 499–513.

Fellbaum, C. (1998). A semantic network of English: The mother of all WordNets. In *EuroWordNet: A Multilingual Database with Lexical Semantic Networks* (pp. 137–148). Dordrecht: Springer.

Fourtassi, A., Bian, Y., & Frank, M. C. (2020). The growth of children's semantic and phonological networks: Insight from 10 languages. *Cognitive Science*, *44*(7), e12847.

Fox, M. D., Zhang, D., Snyder, A. Z., & Raichle, M. E. (2009). The global signal and observed anticorrelated resting state brain networks. *Journal of Neurophysiology*, *101*(6), 3270–3283.

Franklin, N. T., Norman, K. A., Ranganath, C., Zacks, J. M., & Gershman, S. J. (2020). Structured event memory: A neuro-symbolic model of event cognition. *Psychological Review*, *127*(3), 327–361.

Griffiths, T. L., Steyvers, M., & Firl, A. (2007). Google and the mind: Predicting fluency with PageRank. *Psychological Science*, *18*(12), 1069–1076.

Griffiths, T. L., Steyvers, M., & Tenenbaum, J. B. (2007). Topics in semantic representation. *Psychological Review*, *114*(2), 211.

Grohe, M. (2020, June). word2vec, node2vec, graph2vec, x2vec: Towards a theory of vector embeddings of structured data. In *Proceedings of the 39th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems* (1–16).

Grover, A., & Leskovec, J. (2016, August). node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 855–864).

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.

Günther, F., Rinaldi, L., & Marelli, M. (2019). Vector-space models of semantic representation from a cognitive perspective: A discussion of common misconceptions. *Perspectives on Psychological Science*, *14*(6), 1006–1033.

Hamilton, W., Ying, Z., & Leskovec, J. (2017). Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems* (pp. 1024–1034).

Hills, T. T., Kalff, C., & Wiener, J. M. (2013). Adaptive Lévy processes and area-restricted search in human foraging. *PLoS One*, *8*(4), e60488.

Hills, T. T., Maouene, M., Maouene, J., Sheya, A., & Smith, L. (2009). Longitudinal analysis of early semantic networks: Preferential attachment or preferential acquisition? *Psychological Science*, *20*(6), 729–739.

Hills, T. T., Jones, M. N., & Todd, P. M. (2012). Optimal foraging in semantic memory. *Psychological Review*, *119*(2), 431.

Hills, T. T., Todd, P. M., & Jones, M. N. (2015). Foraging in semantic fields: How we search through memory. *Topics in Cognitive Science*, *7*(3), 513–534.

Howell, S. R., Jankowicz, D., & Becker, S. (2005). A model of grounded language acquisition: Sensorimotor features improve lexical and grammatical learning. *Journal of Memory and Language*, *53*(2), 258–276.

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.

Jacoby, L. L., Jones, T. C., & Dolan, P. O. (1998). Two effects of repetition: Support for a dual-process model of know judgments and exclusion errors. *Psychonomic Bulletin & Review*, *5*(4), 705–709.

Jameel, M., Bouraoui, Z., & Schockaert, S. (2018). Unsupervised learning of distributional relation vectors. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics* (Long Papers), 23–33, Melbourne, Australia.

Jamieson, K. G., Jain, L., Fernandez, C., Glattard, N. J., & Nowak, R. D. (2015, January). NEXT: A system for real-world development, evaluation, and application of active learning. In *NIPS* (pp. 2656–2664).

Jones, M. N., Hills, T. T., & Todd, P. M. (2015). Hidden processes in structural representations: A reply to Abbott, Austerweil, and Griffiths. *Psychological Review*, *122*(3), 570–574. https://doi.org/10.1037/a0039248

Jones, M. N., & Mewhort, D. J. (2007). Representing word meaning and order information in a composite holographic lexicon. *Psychological Review*, *114*(1), 1.

Jones, M., & Recchia, G. (2010). You can't wear a coat rack: A binding framework to avoid illusory feature migrations in perceptually grounded semantic models. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 32, No. 32).

Jones, M. N., Willits, J., Dennis, S., & Jones, M. (2015). Models of semantic memory. In J. T. Townsend & J. R. Busemeyer (Eds.), *Oxford handbook of mathematical and computational psychology* (pp. 232–254). New York, NY: Oxford University Press.

Kenett, Y. N. (2019). What can quantitative measures of semantic distance tell us about creativity? *Current Opinion in Behavioral Sciences*, *27*, 11–16.

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., & Austerweil, J. L. (2016). Examining search processes in low and high creative individuals with random walks. In A. Papafragou, D. Grodner, D. Mirman, & J. C. Trueswell (Eds.), *Proceedings of the 38th Annual Meeting of the Cognitive Science Society* (pp. 313–318). Austin, TX: Cognitive Science Society.

Kenett, Y. N., & Faust, M. (2019). A semantic network cartography of the creative mind. *Trends in Cognitive Sciences*, *23*(4), 271–274.

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), 1470–1489.

Kenett, Y. N., Levy, O., Kenett, D. Y., Stanley, H. E., Faust, M., & Havlin, S. (2018). Flexibility of thought in high creative individuals represented by percolation analysis. *Proceedings of the National Academy of Sciences*, *115*(5), 867–872.

Kenett, Y. N., & Thompson-Schill, S. L. (2020). *Novel conceptual combination can dynamically reconfigure semantic memory networks*. PsyArXiv.

Kumar, A. A. (2021a). *Modeling semantic structure and spreading activation in retrieval tasks* [Doctoral dissertation]. Washington University in St Louis. ProQuest Dissertations Publishing. https://www.proquest.com/docview/2520826914

Kumar, A. A. (2021b). Semantic memory: A review of methods, models, and current challenges. *Psychonomic Bulletin & Review*, *28*, 40–80.

Kumar, A. A., Balota, D. A., & Steyvers, M. (2020). Distant connectivity and multiple-step priming in large-scale semantic networks. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *46*(12), 2261–2276. https://doi.org/10.1037/xlm0000793

Kumar, A. A., Steyvers, M., & Balota, D. A. (under review). Semantic memory search and retrieval in a cooperative word game: A comparison of associative and distributional semantic models.

Lorch, Jr, R. F. (1982). Priming and search processes in semantic memory: A test of three models of spreading activation. *Journal of Verbal Learning and Verbal Behavior*, *21*(4), 468–492.

Marelli, M., Gagné, C. L., & Spalding, T. L. (2017). Compounding as abstract operation in semantic space: Investigating relational effects through a large-scale, data-driven computational model. *Cognition*, *166*, 207–224.

Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. Cambridge, MA: MIT Press.

McCloskey, M., & Glucksberg, S. (1979). Decision processes in verifying category membership statements: Implications for models of semantic memory. *Cognitive Psychology*, *11*(1), 1–37.

McLeod, P., Shallice, T., & Plaut, D. C. (2000). Attractor dynamics in word recognition: converging evidence from errors by normal subjects, dyslexic patients and a connectionist model. *Cognition*, *74*(1), 91–114.

McNally, R. J., Robinaugh, D. J., Wu, G. W., Wang, L., Deserno, M. K., & Borsboom, D. (2015). Mental disorders as causal systems: A network approach to posttraumatic stress disorder. *Clinical Psychological Science*, *3*(6), 836–849.

McNamara, T. P. (2005). *Semantic priming: Perspectives from memory and word recognition*. New York, NY: Psychology Press.

McRae, K. (2004). Semantic memory: Some insights from feature-based connectionist attractor networks. *The Psychology of Learning and Motivation: Advances in Research and Theory*, *45*, 41–86.

McRae, K., De Sa, V. R., & Seidenberg, M. S. (1997). On the nature and scope of featural representations of word meaning. *Journal of Experimental Psychology: General*, *126*(2), 99–130.

Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). *Efficient estimation of word representations in vector space*. arXiv preprint arXiv:1301.3781.

Miller, G. A. (1995). WordNet: A lexical database for English. *Communications of the ACM*, *38*(11), 39–41.

Mitchell, J., & Lapata, M. (2010). Composition in distributional models of semantics. *Cognitive Science*, *34*(8), 1388–1429.

Musz, E., & Thompson-Schill, S. L. (2015). Semantic variability predicts neural variability of object concepts. *Neuropsychologia*, *76*, 41–51.

Musz, E., & Thompson-Schill, S. L. (2017). Tracking competition and cognitive control during language comprehension with multi-voxel pattern analysis. *Brain and Language*, *165*, 21–32.

Neely, J. H. (1977). Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited-capacity attention. *Journal of Experimental Psychology: General*, *106*(3), 226–254.

Neely, J. H. (1991). Semantic priming effects in visual word recognition: A selective review of current findings and theories. *Basic Processes in Reading: Visual Word Recognition*, *11*(1), 264–336.

Nelson, D. L., Dyrdal, G. M., & Goodmon, L. B. (2005). What is preexisting strength? Predicting free association probabilities, similarity ratings, and cued recall probabilities. *Psychonomic Bulletin & Review*, *12*(4), 711–719.

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.

Osgood, C. E. (1952). The nature and measurement of meaning. *Psychological Bulletin*, *49*(3), 197.

Pennington, J., Socher, R., & Manning, C. (2014). Glove: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)* (pp. 1532–1543).

Perozzi, B., Al-Rfou, R., & Skiena, S. (2014, August). Deepwalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 701–710).

Quillian, M. R. (1967). Word concepts: A theory and simulation of some basic semantic capabilities. *Behavioral Science*, *12*(5), 410–430.

Richie, R., Aka, A., & Bhatia, S. (in preparation). Free association in bidirectional memory networks.

Roget, P. M. (1911). Roget's thesaurus of English words and phrases. https://www.gutenberg.org/ebooks/10681

Rumelhart, D. E., Hinton, G. E., & McClelland, J. L. (1986). A general framework for parallel distributed processing. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, *1*(45–76), 26.

Siew, C. S. (2019). spreadr: An R package to simulate spreading activation in a network. *Behavior Research Methods*, *51*(2), 910–929.

Siew, C. S., Wulff, D. U., Beckage, N. M., & Kenett, Y. N. (2019). Cognitive network science: A review of research on cognition through the lens of network representations, processes, and dynamics. *Complexity*, *2019*.

Singh, P. (2002). The open mind common sense project.

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), 214–241.

Smith, E. R., & DeCoster, J. (2000). Dual-process models in social and cognitive psychology: Conceptual integration and links to underlying memory systems. *Personality and Social Psychology Review*, *4*(2), 108–131.

Smith, K. A., Huber, D. E., & Vul, E. (2013). Multiply-constrained semantic search in the Remote Associates Test. *Cognition*, *128*(1), 64–75.

Solomon, S. H., Medaglia, J. D., & Thompson-Schill, S. L. (2019). Implementing a concept network model. *Behavior Research Methods*, *51*(4), 1717–1736.

Speer, R., Chin, J., & Havasi, C. (2016). *Conceptnet 5.5: An open multilingual graph of general knowledge*. arXiv preprint arXiv:1612.03975.

Stella, M., Beckage, N. M., & Brede, M. (2017). Multiplex lexical networks reveal patterns in early word acquisition in children. *Scientific Reports*, *7*, 46730.

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.

Steyvers, M., Shiffrin, R. M., & Nelson, D. L. (2005). Word association spaces for predicting semantic similarity effects in episodic memory. In A. F. Healy (Ed.), *Decade of behavior. Experimental cognitive psychology and its applications* (pp. 237–249). American Psychological Association. https://doi.org/10.1037/10895-018

Tang, J., Qu, M., Wang, M., Zhang, M., Yan, J., & Mei, Q. (2015, May). Line: Large-scale information network embedding. In *Proceedings of the 24th International Conference on World Wide Web* (pp. 1067–1077).

Treisman, A. M. (1969). Strategies and models of selective attention. *Psychological Review*, *76*(3), 282–299.

Vankrunkelsven, H., Verheyen, S., Storms, G., & De Deyne, S. (2018). Predicting lexical norms: A comparison between a word association model and text-based word co-occurrence models. *Journal of Cognition*, *1*(1), 45.

Vitevitch, M. S., Ercal, G., & Adagarla, B. (2011). Simulating retrieval from a highly clustered network: Implications for spoken word recognition. *Frontiers in Psychology*, *2*, 369.

Wulff, D. U., De Deyne, S., Jones, M. N., Mata, R., & Aging Lexicon Consortium. (2019) New perspectives on the aging lexicon. *Trends in Cognitive Sciences*, *23*(8), 686–698.

Zemla, J. C., & Austerweil, J. L. (2018). Estimating semantic networks of groups and individuals from fluency data. *Computational Brain & Behavior*, *1*(1), 36–58.