

Manuscript version: Author's Accepted Manuscript

The version presented in WRAP is the author's accepted manuscript and may differ from the published version or Version of Record.

Persistent WRAP URL:

<http://wrap.warwick.ac.uk/157849>

How to cite:

Please refer to published version for the most recent bibliographic citation information. If a published version is known of, the repository item page linked to above, will contain details on accessing it.

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk.

Is the mind a network? Maps, vehicles, and skyhooks in cognitive network science

Thomas T. Hills¹ & Yoed N. Kenett²

¹Department of Psychology, University of Warwick, Coventry CV4 7A, United Kingdom

²Faculty of Industrial Engineering and Management, Technion – Israel Institute of Technology, Haifa
3200003 Israel

Corresponding author:

Yoed N. Kenett

Technion – Israel Institute of Technology

Faculty of industrial Engineering and Management

Haifa, Israel 3200003

Email: yoedk@technion.ac.il

Phone number: +972-4-829-4430

Abstract

Cognitive researchers often carve cognition up into structures and processes. Cognitive processes operate on structures, like vehicles driving over a map. Language alongside semantic and episodic memory are proposed to have structure, as are perceptual systems. Over these structures, processes operate to construct memory and solve problems by retrieving and manipulating information. Network science offers an approach to representing cognitive structures and has made tremendous inroads into understanding the nature of cognitive structure and process. But is the mind a network? If so, what kind? In this article, we briefly review the main metaphors, assumptions, and pitfalls prevalent in cognitive network science (maps and vehicles; one network/process to rule them all), highlight the need for new metaphors that elaborate on the map-and-vehicle framework (wormholes, skyhooks, and generators), and present open questions in studying the mind as a network (the challenge of capturing network change, what should the edges of cognitive networks be made of, and aggregated vs. individual-based networks). One critical lesson of this exercise is that the richness of the mind as network approach is a powerful tool in its own right; it has helped to make our assumptions more visible, generating new and fascinating questions, and enriching the prospects for future research. A second lesson is that the mind as network--though useful--is incomplete. The mind is not a network, but it may contain them.

Keywords: Cognitive networks, representation, process, language, memory, cognition

1. Introduction

In recent years, cognitive networks have received tremendous interest (Baronchelli, Ferrer-i-Cancho, Pastor-Satorras, Chater, & Christiansen, 2013; Borge-Holthoefer & Arenas, 2010; Castro & Siew, 2020; Karuza, Kahn, Thompson-Schill, & Bassett, 2017; Siew, Wulff, Beckage, & Kenett, 2019). Cognitive networks refer to the application of methods from network science to study the complexity of cognitive systems, mainly language and memory (Castro & Siew, 2020; Siew et al., 2019). The success of cognitive networks is due to a number of factors, including the ability to quantify structural characteristics, such as nearness and farness, between behaviors, concepts, and memories (Kenett, Levi, Anaki, & Faust, 2017; Kumar, Balota, & Steyvers, 2020), the potential for enrichment or degradation of the network map—such as in creativity (Kenett, 2018, 2019; Kenett & Faust, 2019b) and age-related cognitive decline (Borge-Holthoefer, Moreno, & Arenas, 2011; Kenett et al., 2018)—and the possibility that aspects of cognitive control (executive function) influence network navigation (Hills, Mata, Wilke, & Samanez-Larkin, 2013).

As two avid users of network science, we constantly discuss and often disagree about big details. For example, is creativity explained by a change in the map (i.e., the network representation; Kenett & Faust, 2019b) or is it a consequence of the vehicle (i.e., the cognitive processes such as cognitive control that navigate the network; Silvia, 2015)? Is memory one representation or is it many? If it is many, then memory search may be amenable to ‘short-cuts’ or ‘wormholes’, by which travel between two distant places in a network can be shortened by briefly traveling via another representation (Wulff, Hills, & Hertwig, 2020). How do we separate the map from the vehicle experimentally? If one researcher infers a difference in representations between two groups (e.g., young vs. old), might that difference also be explained by a change in the processes, such as cognitive control, without a difference in the underlying networks (Siew et al., 2019)?

Questions like these are often the outcome of better understanding the metaphors we have adopted and the assumptions we have made. Cognitive network science is a technical *and* theoretical framework. As Jones and Love (2011; p. 170) describe, there is “a danger of confusing technical advances with theoretical progress, and the allure of the former can lead to the neglect of the latter. As the new framework develops, it is critical to keep the research tied to

certain basic questions, such as: What theoretical issues are at stake? What are the core assumptions of the approach?”

Our goal here is to take these concerns seriously in an honest self-appraisal of our discipline and its theoretical contributions to understanding the mind. In what follows, we will discuss the main metaphor that dominates cognitive network science—the map and vehicle model—and the assumptions it often elicits. Then we will highlight alternative metaphors (wormholes, generators, and skyhooks), based on varied empirical findings. Finally, we will highlight core challenges that cognitive network science must address to advance our understanding of the complexity of the human mind.

2. The map and vehicle framework

William James presented a metaphor of mental ‘travel’ when he proposed that we search memory like we search for a lost item in our house, first checking one room and then another (James, 1890). This invites one to imagine our memories as a kind of map—or as Tolman put it, a *cognitive map*—over which some kind of mental vehicle travels (Tolman & Tolman, 2004). Indeed, any account of cognitive phenomenon requires a description of both what the cognitive system knows (the representation) and what it does (the process) with that information (Anderson, 1990; Estes, 1975).

Take, for example, the case of modelling the semantic fluency task, which is often used in clinical practice (Kenett & Faust, 2019a): “say all the animals you can think of.” Here the researcher’s goal is to be able to predict the list of animal names that an individual is likely to produce: e.g., *dog, cat, monkey, giraffe, etc.* To do this, one first needs to represent how animals are related to one another in the mind. A mathematical description of this space, representing the distances between different concepts, works nicely to represent the map. Such a matrix (or graph) is a network. Second, one needs to formally describe some kind of process that acts over that network, allowing names to be activated in a series (Hills, Jones, & Todd, 2012; Zemla, Kenett, Jun, & Austerweil, 2016). The representation is a map, and the process is a vehicle; together they make explicit the cognitive components necessary to understand how the mind solves this task.

This *map and vehicle framework* represents a core assumption that underlies many network science approaches to cognition. The instantiation of the map and vehicle model varies across researchers; sometimes there are multiple maps and sometimes there are many vehicles, but the conceptual core is the same. Below, we use this core framework to unpack a number of potential assumptions about the practice of cognitive network science, while also highlighting some of the open questions that remain to be addressed and exploring what it would mean for the mind to be a network.

2.1. One Network to Rule Them All

Suppose you assume that the network on which you model your semantic fluency data is based on the same representation for everyone (**Figure 1**). So, no matter how old or young, or how much a lover of Attenborough or hater of the outdoors, everyone is assumed to know the same animals and to represent them in the same way. We call this approach the *one-network-to-rule-them-all assumption*. Someone trying to infer the process people use to navigate this network will find that nature lovers have excellent memory search processes. The home-bodies, on the other hand, will appear to be mentally impaired, unable to retrieve information from the network that everyone is assumed to have. This may sound silly, but exactly the same approach has been used by Hills, Mata, et al. (2013) to make inferences about executive cognitive control processes in the aging mind. *If you assume everyone has the same network, then you guarantee the differences will be in the process.*

2.2. One Process to Rule Them All

Alternatively, one can imagine that everyone has the same process and look for differences in the representation (**Figure 1**). This is the *one-process-to-rule-them-all assumption*. For example, you might ask people to produce associations for a set of words. This data—producing targets in response to cues—can be used to construct networks (De Deyne, Kenett, Anaki, Faust, & Navarro, 2016; De Deyne, Navarro, Perfors, Brysbaert, & Storms, 2019). If the participants are split into different groups based on some other measurement, such as intelligence (low vs. high), the networks constructed out of their aggregate data might look different. One might then conclude that low- and high-intelligent individuals represent information differently. Again, this may sound silly, but exactly the same approach was used by Kenett et al. to evaluate

the semantic representations of low and high creative individuals (Kenett & Faust, 2019b). *If you assume everyone has the same process, then you guarantee the differences will be in the representation.*

2.3. Can representation and process be disentangled?

The above examples are necessarily simplified, but things can get tricky in a hurry. A specific example (described in more detail by Jones, Hills, & Todd, 2015), involves two different approaches to modelling the semantic fluency task described above. Both took a map and vehicle approach, and both used a one-network-to-rule-them-all assumption. One constructed the network out of free associations, derived from another database of free associations wherein people said the first words that came to mind when presented with a cue word, like *cat* (Abbott, Austerweil, & Griffiths, 2015). The second constructed the network using a semantic space model applied to a corpus natural language, based on using patterns of word-occurrences to detect word similarity (Hills et al., 2012). For search over the free association network, Abbott et al. (2015) assumed a process (i.e., vehicle) made up of random walkers, who moved from the last word produced randomly over the network to activate nearby words. For the semantic space network, Hills et al. (2012) assumed the process was a common probabilistic choice rule (the Luce choice rule) that chose nearby words in proportion to their similarity with the previously recalled word. In this second case, if there were no nearby words, a secondary ‘switching’ process was assumed to choose a new word based on its frequency in natural language—leaving the local constraints of the network to jump to a new location. Abbott et al. (2015) demonstrated that the two approaches were able to generate similar phenomenology using different processes (random walkers versus a switching process) and different representations (free associations versus semantic space). Following Anderson (1978), Abbott et al. (2015) put it like this: “behavior that seems like the signature of one mechanism can sometimes be produced by others.”

In relation to the core question of this article, if the mind is a network, the natural extension to this question is “which one”? Theoretical progress in cognitive network science requires that we become more proficient at comparing our alternatives. There are many ways to do this: model competitions, cross-validation, model recovery, parameter evaluations, qualitative comparisons to various datasets, and so on (e.g., Shiffrin, Lee, Kim, & Wagenmakers, 2008).

With respect to representation and process, we can independently sample cognitive control measures (Hills et al., 2013; Hills & Pachur, 2012; Kenett, Beaty, Silvia, Anaki, & Faust, 2016), introduce secondary tasks that interfere with control processing (Rosen & Engle, 1997), use neuroimaging to identify control processing (Beaty, Benedek, Silvia, & Schacter, 2016), and evaluate more nuanced differences in the behaviour, such as reaction time data (Kenett et al., 2017; Kumar et al., 2020). Given a sufficiently large set of possible processes and representations, we are likely to find a variety of process-representation pairs that are difficult to discriminate without formal model comparisons.

Though it is debated under what conditions we can investigate process and representation independently (Siew et al., 2019), the reality is that if we cannot discriminate representation and process our ability to make general inferences from cognitive network research is highly constrained and likely to be presumptive whenever we attribute explanations to structure or process without testing the alternatives.




Figure 1: **The core metaphor of cognitive network science and two typical assumptions.** Map and vehicle. This is a common implicit framework for many network approaches to cognition in which nodes represent behavioural outputs. The network is the structural representation over which some process (e.g., vehicle) navigates. **One-process-to-rule-them-all** assumes that all participants have the same process and that only the structure of the network representation may change. **One-network-to-rule-them-all** assumes that everyone has the same representation and one can therefore examine individual differences in processes (e.g., a Luce choice rule, random walkers, a sailboat).

3. Alternative cognitive network metaphors

The observations we discuss above are meant to shine light on one overarching assumption (the map and vehicle) and two common practical assumptions (one-process/network-to-rule-them-all) that may often be at work when taking a network approach to cognition. But let us take a second to challenge the most alluring assumption of all: is the mind a network?

If the mind is a network, the map-and-vehicle approach already suggests that a network (or map) is incomplete. From a purist perspective, a network and the mathematical matrix that represents it does not *do* anything; one must apply a secondary process to deliver some output in response to some input. It is instead the knowledge representation of the system that may be mathematically equivalent to a network. Metaphors of knowledge representations have ranged across such diverse things as a cow's stomach (Hintzman, 1974), an acid bath (Posner & Konick, 1966), a library card catalogue (Forster, 1978), steam engines, and a digital computer (for a review, see Turvey & Moreno, 2006). Still, the most persistent metaphor is of memory as akin to a physical space (Roediger, 1980), satisfying the network assumption of a quantifiable distance between 'places' in the space. Plato and Aristotle (cited in Roediger, 1980) considered memory as a kind of wax tablet. Freud (1920) revived this idea with his mystic writing pad. James (1890) proposed that memory was like a house one could walk around in and Collins and Quillian (1969) expanded the idea to a subway map. Collins and Loftus further developed this idea to suggest that memory is structured as a network that constrains a process of spreading activation (Collins & Loftus, 1975). Hills (2006) went on to suggest that the similarity between external and internal representations facilitated the evolution of a common search process to navigate them.

This leads us to the map and vehicle framework, which has had substantial success in describing cognitive phenomena, as cited throughout this article. Graphs (i.e., matrices) and their corresponding networks quantify the structure of cognitive space and provide the kind of pin-headed angel counting required to distinguish one theory from another. Nevertheless, the map and vehicle metaphor is but one way to achieve that. What are the alternatives?

3.1. Wormholes

A common assumption in cognitive network science is that the mind uses a single fixed network. Even if different people have different network representations, they each have but one. One cannot jump from one network to another within the same mind. The network one is assumed to have is also rigid and inflexible—one cannot emphasize certain edge properties over others, as edges are often deemed to have a binary value (present or absent) or to be a weighted output of their subcomponents that are themselves no longer separable.

However, consider searching for country names from memory (“Switzerland”, “Germany”, “The UK”, etc). In this case, we may sometimes feel we are searching based on one kind of feature—such as the spatial proximity of a shared border—and at other times feel we are searching based on phonological similarity, as one might have when recalling the ‘Stans’, i.e., “Kazakhstan”, “Kyrgyzstan”, “Pakistan”, and so on. The ability to search via different features may suggest the existence of more than one localist network representation and the capacity to switch between them, or similarly, the capacity to restructure a single network by emphasizing or de-emphasizing the different kinds of features that make up its edges. We consider these cognitive wormholes, because they would allow cognition to bend internal space in a directed fashion by emphasizing one network representation to another (**Figure 2**).

Wulff, Hills, and Hertwig (2020) studied a milder wormhole hypothesis by having people retrieve names of countries either by continent or by first letter (from A to Z). They hypothesized that a continent-based search would emphasize a spatial representation, as if the person were imagining a map. Similarly, they hypothesized that an alphabetical search would emphasize a phonological representation, as if the person were sounding out the first letter. The data suggested that people followed these predictions. However, evidence of phonological retrieval was apparent in the by-continent condition, and evidence of spatial retrieval was apparent in the alphabet condition. This suggests that people can bend their internal representations by emphasizing different edge properties (e.g., phonology or spatial proximity) in the underlying network—even though they do not appear to be able to separate these in entirety.

This inability to completely separate networks is consistent with tip of the tongue states, in which people express difficulties in retrieving names and get stuck on misleading phonemes, even though they can also access semantic information (Brown & McNeill, 1966; Vitevitch, Chan, & Goldstein, 2014). Thus, while some evidence of wormhole space bending exists, it seems less clear that memory is represented by separable networks.

The above data suggest that task context can influence the structure of a searchable cognitive representation, but that it is likely to remain a single cognitive representation (Kumar, 2021), though it appears to be a little bendy. To get from one place to another, one must search its structure locally or, as noted in the previous Section 2.2, use control processes that escape the local structural constraints. **Importantly, it illustrates that there is an interaction between the process and structure. Therefore, a one network/process to rule them all model is inadequate,**

challenging us to consider different frameworks that allow us to advance our thinking on the relation between them.

3.2. Skyhooks

Can a one network to rule them all model account for human performance in cognitive tasks? As we discuss in section 2.3., two competing models predict performance in a semantic fluency task: random walks and clustering-switching processes. Importantly, a complete model of the task must account for controlled, goal-directed cognition, which may require an integrated approach that combines elements of both (see Goñi et al., 2010 for one such example).

As discussed, one popular model suggests that when participants are performing a memory search task they switch between a ‘local’ representation (the network) and a ‘global’ frequency list (e.g., Hills et al., 2012; Hills et al., 2013; Hills & Pachur, 2012). This capacity to jump using frequency is somewhat similar to a wormhole, except that frequency has no local structure. Having just found the 100th most frequent animal does not—as far as we know—make it easier to retrieve the 99th most frequent animal. In this sense, frequency does not shrink the distance between targets in memory, it simply allows one to escape resource-poor locations on the internal map, a process we suggest is similar to a *skyhook* (**Figure 2**).

Relevant to our discussion on representation and process, these skyhook-like global transitions are correlated with measures of executive cognitive control (Hills et al., 2013; Hills & Pachur, 2012; Rosen & Engle, 1997). Executive cognitive control is also dynamically involved in creative thinking (Chrysikou, 2019). This suggests a nice analogy between thinking out of the box and getting off your cognitive map. Individuals with better performance on measures of executive control retrieve more items from memory and are better at retrieving original ideas (Beatty, Kenett, Hass, & Schacter, 2019; Volle, 2018). Thus, executive control can potentially evaluate the local prospects in the representation and act as a skyhook to transfer attention elsewhere if local opportunities are looking sparse (**Figure 2**).

The skyhook metaphor adds an additional process to the “is the mind a network” question, which currently stands at map plus vehicle. This new skyhook process is related to effortful cognitive control because it is constrained by a cognitive load. It instantiates the effortful component in psychology’s well-established distinctions between automatic and

effortful processing (Kahneman, 2011). Correspondingly, more automatic processes may encompass aspects of the vehicle when it is on the map.

3.3. Generators

Another common observation about cognitive representations is that they change over time. This requires us to invoke processes associated with network change, such as learning and forgetting. New nodes and edges can be formed and potentially lost; we can learn new object concepts and their associations. We can also combine existing nodes to create new nodes, what one might call cognitive griffins (a griffin had the head and wings of an eagle and the body of a lion). Indeed, we know that memory works like this; Bartlett's (1995) work in *Remembering* and much research since has shown that much of cognition is constructive. Stories evolve in their retellings (Jagiello & Hills, 2018; Moussaïd, Brighton, & Gaissmaier, 2015) and we can construct alternatives that we have never experienced (Hills, 2019). Thus, nodes can in principle form and move in relation to other nodes. The map itself can change, and any metaphor of the mind must be able to account for such change.

Though network learning models exist for evaluating how concepts might be learned from their environment (Engelthaler & Hills, 2019; Hills, Maouene, Maouene, Sheya, & Smith, 2009a), these models largely aim to predict what order, for example, children will acquire new words. A deeper question asks how concepts are formed and may change in relation to one another. More specifically, what are nodes and edges constructed of? Understanding the processes that generate nodes may help us understand how they can change.

Neural networks may offer a useful process metaphor for understanding node behavior. Neural networks are different kinds of networks than cognitive networks, both in relation to structure and process. Structurally, cognitive networks usually adopt a localist representation—*cat* is a single node—whereas neural networks adopt a distributed representation—*cat* is distributed across multiple nodes. In terms of process, how nodes change is not yet well-described by cognitive network science. However, neural networks learn using prediction error and alter their distributed representations via feedback (Lecun, Bengio, & Hinton, 2015).

A particularly useful instantiation of a neural network that may be analogous with neural processes in the brain (Gershman, 2019) is a *generative adversarial network* (Figure 2). A generative adversarial network contains a generator and a discriminator: the generator learns to

generate data (e.g., cats) which can fool the discriminator which learns to detect data categories. Both the generator and the discriminator learn through their interaction. A cognitive instantiation of a generative adversarial network can both learn to infer structure (e.g., categories) from data, but also produce data from categories, producing outputs that have never been seen before (e.g., novel faces). Note that the discriminator classifies the data into specific categories (our classic idea of nodes), but can also produce outputs that lie between other outputs, merging details from multiple starting categories (Karras, Laine, & Aila, 2019).

A cognitive system with similar abilities could create novel nodes, such as the meaning of *cabbage cage* or *gravy orgy*, which also represent the kind of novel high-entropy stimuli associated with humor (Siew, Engelthaler, & Hills, under review). Moreover, if as network scientists we want to imagine that the nodes in our networks represent constellations of neural activity in a brain, then we must accept that *cabbages*, *cages*, and *cabbage cages* all represent different patterns of activity across a distributed representation (e.g., Musz & Thompson-Schill, 2015). There is therefore a potential infinity of neural states and associated nodes in our networks.

This truth is embedded in the neuroscience: short- and long-term synaptic plasticity means that constellations of neural activity involve many thousands of continuously tuning synapses, making our ‘network’ nodes moving targets (Holtmaat & Svoboda, 2009). This invites us to think of our output nodes as constructions of other nodes (“wavy”, or some other probabilistic variation, e.g., Tenenbaum, Kemp, Griffiths, & Goodman, 2011), somewhat similar to the interactive activation model of word recognition with adjustable weights (McClelland & Rumelhart, 1981).

A more recent model of semantic representation shares an architecture similar to our generative analogy (Jamieson, Avery, Johns, & Jones, 2018; McKoon & Ratcliff, 1992). In this architecture, words are not abstracted into prototypes by averaging across the contexts in which a word appears—which is the more typical method of distributional semantics. Instead, words preserve their idiosyncratic regularities in an exemplar representation, which captures individual contexts. This allows a retrieval process to recover meanings associated with various homonyms and polysemes (e.g., *break* as in *smash* or *break* as in *report*)—a standard problem for distributional semantic models. The model proposed by Jamieson et al. (2018)—the instance theory of semantics—distributes concept information across features and allows memory traces

to be constructed at retrieval based on differential feature activation. This may offer a process based mechanism for deriving wormhole-like behaviour and it also has the properties of the generator we propose here.

4. Core challenges in cognitive network science

Wormholes, skyhooks, and generators offer metaphors to further elucidate the underlying assumptions of the map and vehicle framework by suggesting how they might be otherwise. There are several additional core challenges that cognitive network scientists should be aware of, which represent both pitfalls and opportunities: What “stuff” make up the edges in cognitive networks, how can we measure the effect of change and context in cognitive networks, and the ability to measure individuals, not just be aggregated data of groups.




Figure 2. Alternative to the standard map and vehicle framework. **Wormholes** allow individuals to take shortcuts across a network by using a separate network that contains a more direct path to the destination. **Skyhooks** represent cognitive processes that cause motion on the network independent of the network structure. **Generators** represent multipartite networks in which the top-level localist network is constructed from activations of a lower-level distributed network.

4.1. Network Chimera

What holds the mind’s nodes together? What are its edges (the links between nodes)? Importantly, even in the related field network neuroscience, researchers still discuss what the edges should be in functional MRI connectivity-based networks (Lurie et al., 2019). In cognitive network research, the discussion is a bit more muted, but needs to take place. Researchers have wielded a strikingly large set of materials with which to link their network nodes, such as phonological similarity (FAT and CAT are neighbors; Vitevitch, 2008; Vitevitch & Castro, 2015), perceptual features (a banana is yellow and so is a school bus, Sizemore, Karuza, Giusti, & Bassett, 2017), functional features (a hammer hits and so does a bat; Engelthaler & Hills,

2019), co-occurrences in language (Beckage, Smith, & Hills, 2011), semantic similarity derived from text (Bhatia, 2017; Hills et al., 2010), orthographic similarity (Siew, 2018), similarity of responses to items in questionnaires (Christensen, Kenett, Aste, Silvia, & Kwapil, 2018), and many other things (Siew et al., 2019). Is one of these *the* answer?

It is tempting to suggest that edges in a memory network, for example, can be inferred from free association data (Abbott et al., 2015; De Deyne et al., 2019; Kenett et al., 2017; Kumar, 2021; Kumar et al., 2020). However, the cognitive mechanisms underlying free association and its ability to reliably represent individual-based and group-based representations are far from understood (Nelson, McEvoy, & Dennis, 2000). We can measure associations from people by giving them a *cue* word (e.g., cat) and asking for the first word that comes to mind, the *target* (Nelson, McEvoy, & Schreiber, 2004). Other free association data collection approaches require participants to generate multiple responses (De Deyne et al., 2019; Kenett, Kenett, Ben-Jacob, & Faust, 2011). There are large collections of these free association norms (e.g., De Deyne et al., 2019; Nelson et al., 2004). Free associations are a common approach to estimate networks and have been used to model everything from memory retrieval (Steyvers & Tenenbaum, 2005) to word learning (Hills, Maouene, Maouene, Sheya, & Smith, 2009b).

However, Jones et al. (2015) argued that free associations have a *Turk problem*. The problem is named after the 18th century chess-playing automaton, the Mechanical Turk, which secretly housed a chess-playing boy. Similar to the Turk, free associations (especially when multiple response are required) may secretly house one or more processes. In other words, they may not be a process-free read out of an individual's underlying cognitive representation. If researchers are trying to explain the cognitive processes that produce a certain behavior, using free associations may obscure that process. Alternatively, if one is representation-process agnostic—and indifferent to the underlying mechanism driving the behavior—then free associations may be easy and ideal. If nothing else, they are often a good place to start in identifying differences, before one moves on to exploring the processes that give rise to those differences (see Hills, Maouene, Riordan, & Smith, 2010 for an example of exploring what may drive a free association pattern in a network).

What edges work best are likely to depend on the question one is trying to answer. For children learning early words, co-occurrences in language appear to outcompete other measures (including free associations) in predicting order of word learning (Hills et al., 2010; Hills et al.,

2009b). But different questions about word learning provide roles for different kinds of networks (Hills & Siew, 2018). In some cases, evidence against one kind of network representation (e.g., features: being furry, having eyes, etc) is revived when the process used to construct edges is revised. For example, Hills et al. (2009b) showed that feature networks were not effective at predicting the order of children's word acquisition when network edges represented one or more shared features. However, Engelthaler and Hills (2017) found that if they reframed feature networks based on distinctiveness—how dissimilar are two objects—they *could* predict word order using features. This invites us to consider that network edges are themselves the outcomes of cognitive processes.

An alternative to committing to a single type of edge is to keep all the different edge types and represent them in different layers, in what are called *multiplex networks*. Various approaches to multiplex networks have proven to be quite enlightening (Stella, 2019, 2020; Stella, Beckage, & Brede, 2017; Stella, Beckage, Brede, & De Domenico, 2018; Stella & Brede, 2016; Stella & Kenett, 2019). Moreover, recent work has shown that the relative layers of multiplex networks (e.g., phonology, co-occurrences, and free associations) vary in their predictive power over time in relation to early word learning (Stella, 2019; Stella et al., 2018)—which indicates a general need to attend to individual differences (e.g., age-related differences) when testing process and representation hypotheses. Moreover, multiplex networks may offer a substrate to explain the wormhole findings discussed above and the short-term network change described below (Kenett & Thompson-Schill, 2020).

4.2. Network Change

Given that cognitive representations may change as we learn or age, it is useful to ask whether and how networks change over time (Figure 3). A few studies have investigated the effect of aging on semantic memory, or the aging lexicon (Wulff, De Deyne, Jones, Mata, & Consortium, 2019). A seeming paradox in aging research is that while many cognitive capacities decline over time, people accumulate more semantic knowledge as they age (Park et al., 2002). Aging cognitive network studies report structural properties of the lexical representation that vary across the lifespan (Cosgrove, Kenett, Beaty, & Diaz, 2021; Dubossarsky, De Deyne, & Hills, 2017; Wulff, Hills, & Mata, 2018; Zortea, Menegola, Villavicencio, & Salles, 2014) and

find that concepts in older adults' semantic memory are less well connected and are more segregated (any pair of concepts in the network are “further apart”) than those of younger adults. For example, Dubossarsky et al. (2017) analyzed the network structure of free associations obtained from a cross-sectional sample across the lifespan to estimate semantic networks for groups of young, middle-aged, and older adults. The authors find a U-shape change in semantic memory properties across the lifespan: Semantic memory starts off as sparse, increases in density towards midlife, which is then followed by increasing sparseness in late life (Dubossarsky et al., 2017). Cosgrove, Kenett, Beaty, and Diaz (2021) have recently shown how such structural changes lead to diminished flexibility in these networks. Importantly, such aging differences across the mental lexicon may be related to changes in retrieval processes, which are also known to change across the lifespan (Salthouse, Atkinson, & Berish, 2003).

Short-term change in cognitive networks may also occur (Kenett & Thompson-Schill, 2020; Yee & Thompson-Schill, 2016). Kenett and Thompson-Schill (2020) estimated semantic networks based on free associations, before and after a short-term cognitive manipulation where they primed participants to process concepts via different strategies. These strategies either had participants focus on the dominant features of these concepts, or on relations between these concepts and other concepts (Kenett & Thompson-Schill, 2020). The authors found that processing concepts based on the relations between them led to changes in the post-manipulation semantic network. In addition, the authors show that in a baseline condition—where pre- and post-semantic networks were estimated without the manipulation task—there were no differences between the two networks, suggesting they remained stable.

However, an alternative interpretation to the findings of Kenett and Thompson-Schill (2020) is that the manipulation task affected participants sensitivity to different features of the concepts, which led to different retrieval processes during the post-manipulation free association task, and not to changes at the representational level. Such an account fits with the instance theory of semantic memory described above (Jamieson et al., 2018), and also dynamical attractor network models, which model how an input to the network (such as thinking about a concept) may change the state into which the network settles depending upon the time course of input activation (e.g., conceptual combinations). For example, Lerner et al. has shown how associative thinking and semantic priming can be modelled based on such attractor-based models (Lerner, Bentin, & Shriki, 2012, 2014; Lerner & Shriki, 2014).

To address this alternative perspective, Kenett and Thompson-Schill (2020) conducted several additional analyses of the behavioral data collected to show that the conceptual combination manipulation similarly affected the free associations generated after the manipulation, compared to those generated before, but that the content of retrieved associations was different, according to the manipulation. These findings may provide indirect support for representation-based, and not process-based, change. However, this evidence can also be explained by the instance theory of semantic memory (Jamieson et al., 2018). Thus, whether the representation level was affected at all, and to what extent these short-term effects persist remains an open question.

4.3. Aggregated Networks

The majority of network research is group-based, aggregating data from multiple individuals to compare differences in representation and process. The problem with such a group-based approach relates to a larger more general issue in cognitive research related to problems of making inferences about individuals from aggregated data (Estes, 1956; Myung, Kim, & Pitt, 2000). Consider the studies described above regarding differences in the semantic networks in low and high creative individuals (Kenett & Faust, 2019b). Creativity varies across individuals, yet these cognitive network studies aggregate participants into low and high creative groups. Aggregation of individual data poses potential problems and it is not well understood in cognitive networks science. Thus, where possible, we recommend collecting individual data.

Some studies have collected individual networks from either the individuals themselves or their environments (Benedek et al., 2017; Hills & Pachur, 2012; Hills & Segev, 2014; Morais, Olsson, & Schooler, 2013; Zemla et al., 2016). For example, Morais, Olsson, and Schooler (2013) applied a “snow-ball” sampling method to collect associative networks for individuals, taking the targets produced during one sampling event as cues in subsequent sampling events, a procedure which took 30 to 60 days to complete. Zemla et al. (2018; 2016) combined multiple re-iterations of the semantic fluency task with Bayesian modelling to estimate an individual’s semantic network for the animal category. Finally, others have used relatedness judgment tasks, where participants rate the relatedness of pairs of words, to estimate the semantic network of these words, an approach that has been conducted across several languages (Benedek et al., 2017; Bernard, Kenett, Ovando-Tellez, Benedek, & Volle, 2019; He et al., 2020). Individual data

can be especially useful for teasing apart differences in representation and process. An individual's network can be collected independently of the processes they use to navigate it, which can be further investigated using additional cognitive control tasks, for example, allowing one to infer that working memory performance predicts network perseveration (e.g., Hills & Pachur, 2012).




Figure 3, Core challenges in cognitive network science. **Network chimera** represents the challenge of identifying the proper edges when there are many to choose from. **Network change** represents the challenge of characterizing how networks and processes change over time. **Aggregate networks** represents the challenges of understanding aggregate networks built from individual data, which potentially hide the underlying processes or structure that gave rise to them.

5. Conclusions

Is the mind a network? Our analysis suggests that the network analogy is far from complete. While the mind clearly represents information, and that representation may be usefully approximated as a network, the characteristics of retrieved knowledge as dealt with by cognitive network science cannot be reduced to a network. The reigning assumptions (one network or process to rule them all) already invite a map and vehicle metaphor. Indeed, it is not yet clear how to capture representational data without involving some form of retrieval process. Though we discuss numerous ways to disentangle these, the mind (which we have only defined from a functionalist perspective) would seem to at least involve process *and* representation.

We also grappled with additional assumptions: bendy networks (wormholes), executive/cognitive control processes (skyhooks), and nodes as dynamic and distributed representations (generators). Given the kinds of questions cognitive scientists are likely to ask about cognition, these too would seem to be required for a complete understanding of cognition. In practice, the emphasis on a map or vehicle varies across practitioners of the network arts. The alternatives we enumerate here (and no doubt others) have yet to be examined and many

theoretical issues remain in mid-embrace. Thus, in answer to the question of “where do we go from here,” **Table 1** provides an overview of existing assumptions as well as the new metaphors and challenges that remain open questions for future work.

It is always useful to examine our metaphors and the limitations they may obfuscate as well as the alternatives that may liberate them. We hope to have captured some of that intention here, recognizing that our approach is still likely to be limited. There is of course additional phenomenology that cognitive network science has yet to address and that we did not discuss in detail. These include problems like chunking in memory, procedural knowledge, personality, episodic memory, the difference between experts and novices in strategic control (e.g., in chess), and so on. These remain exciting open challenges for cognitive network science and we are hopeful that the application of cognitive network science will remain a useful tool in contributing to how we understand these questions.

Nonetheless, one clear finding of this exercise is the substantial progress cognitive network science has made on various fronts in a relatively short amount of time, and the numerous ways in which network science has demonstrated the importance of structure and control in explaining cognitive phenomena. We might turn to the brain for some help in guiding us towards a future cognitive metaphor. The controlled semantic cognition framework suggests that semantic cognition involves the interaction between two neural systems: conceptual knowledge and control processes that guide retrieval (Lambon Ralph, Jefferies, Patterson, & Rogers, 2017; Rogers, Patterson, Jefferies, & Lambone-Ralph, 2015). According to this model, the control processes are “implemented within a distributed neural network that interacts with, but is largely separate from, the network for semantic representation” (Lambon Ralph et al., 2017; p. 49). This framework suggests that it should be possible to study representation and processes separately at the cognitive level. More importantly, it highlights the strengths of incorporating findings and insights from brain research into cognitive theory.

Nevertheless, and despite its limitations and required theoretical and methodological growth, we argue that cognitive network science is a powerful quantitative perspective to capture the complexity of cognition. Its central contribution is its ability to understand how representations differ, how they are acquired, and how they engage with other processes to influence behavior. Moreover, the quantitative language of networks allows us to conduct novel research that combines and cuts across different levels of analysis (neural networks, cognitive

networks, social networks), allowing us to ask new research questions that were not possible to ask before. In sum, cognitive network science poses the potential to significantly impact our theories and operationalization of complex cognitive phenomena. Though assumptions abound, in many cases the alternatives to these assumptions represent untapped opportunities to advance new theoretical and empirical cognitive science.

Table 1: Assumptions, alternatives, and core challenges that lie at the heart of cognitive network science. For each of these issues, we highlight the key questions an avid cognitive network scientist should be cognizant of and examine as they set forth in a cognitive network research.

Assumptions on the network model		
Map and vehicle	Cognitive processes consist of information representations and processes that operate on these representations to generate behavior.	
One-network-to-rule-them-all	If you assume everyone has the same network, then you guarantee the differences will be in the processes.	
One-process-to-rule-them-all	If you assume everyone has the same process, then you guarantee the differences will be in the representation.	
Alternative cognitive network metaphors		
Wormholes	Can people access more than one network independently in the same task, or is there one fixed network?	
Generators	Are nodes fixed or might the network be better represented as a multipartite network, for which the location of nodes in the network are influenced by their relationships with nodes in other networks?	
Skyhooks	Is there a non-network process (e.g., executive attention) that influences movement on the representation?	
Core challenges in cognitive network science		
Network chimera	Are you studying a particular kind of edge or a particular kind of behavior? If the latter, are edges one feature or many? Can you compete different edges against one another?	
Network change	Does the network change in a predictable way with age or experience?	
Aggregated networks	Is the network aggregated data, which may obscure underlying structure or processes? Do you have access to individual data?	

References

- Abbott, J. T., Austerweil, J. L., & Griffiths, T. L. (2015). Random walks on semantic networks can resemble optimal foraging. *Psychological Review*, *122*(3), 558-569.
- Anderson, J. R. (1978). Arguments concerning representations for mental imagery. *Psychological Review*, *85*(4), 249-277.
- Anderson, J. R. (1990). *The adaptive character of thought*. Psychology Press.
- Baronchelli, A., Ferrer-i-Cancho, R., Pastor-Satorras, R., Chater, N., & Christiansen, M. H. (2013). Networks in cognitive science. *Trends in Cognitive Sciences*, *17*(7), 348-360.
- Bartlett, F. C. (1995). *Remembering: A study in experimental and social psychology*. Cambridge University Press.
- Beaty, R. E., Benedek, M., Silvia, P. J., & Schacter, D. L. (2016). Creative cognition and brain network dynamics. *Trends in Cognitive Sciences*, *20*(2), 87-95.
- Beaty, R. E., Kenett, Y., Hass, R., & Schacter, D. L. (2019). A Fan effect for creative thought: Semantic richness facilitates idea quantity but constrains idea quality. *arXiv*.
- Beckage, N., Smith, L., & Hills, T. T. (2011). Small worlds and semantic network growth in typical and late talkers. *PLoS ONE*, *6*(5), e19348.
- Benedek, M., Kenett, Y. N., Umdasch, K., Anaki, D., Faust, M., & Neubauer, A. C. (2017). How semantic memory structure and intelligence contribute to creative thought: a network science approach. *Thinking & Reasoning*, *23*(2), 158-183.
- Bernard, M., Kenett, Y. N., Ovando-Tellez, M. P., Benedek, M., & Volle, E. (2019). *Building individual semantic networks and exploring their relationship with creativity*. Paper presented at the The 41st Annual Meeting of the Cognitive Science Society, Montreal QB, Canada.
- Bhatia, S. (2017). Associative judgment and vector space semantics. *Psychological Review*, *124*(1), 1.
- Borge-Holthoefer, J., & Arenas, A. (2010). Semantic networks: Structure and dynamics. *Entropy*, *12*(5), 1264-1302.
- Borge-Holthoefer, J., Moreno, Y., & Arenas, A. (2011). Modeling abnormal priming in Alzheimer's patients with a free association network. *PLoS ONE*, *6*(8), e22651.
- Brown, R., & McNeill, D. (1966). The "tip-of-the-tongue" phenomenon. *Journal of Verbal Learning and Verbal Behavior*, *6*(325-337).
- Castro, N., & Siew, C. S. Q. (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., 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.
- Chrysikou, E. G. (2019). Creativity in and out of (cognitive) control. *Current Opinion in Behavioral Sciences*, *27*, 94-99.
- Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, *82*, 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.
- De Deyne, S., Kenett, Y. N., Anaki, D., Faust, M., & Navarro, D. J. (2016). Large-scale network representations of semantics in the mental lexicon. In M. N. Jones (Ed.), *Big data in cognitive science: From methods to insights* (pp. 174-202). New York, NY: Psychology Press: Taylor & Francis.

- 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.
- Dubossarsky, H., De Deyne, S., & Hills, T. T. (2017). Quantifying the structure of free association networks across the life span. *Developmental Psychology*, 53(8), 1560-1570.
- Engelthaler, T., & Hills, T. T. (2017). Feature biases in early word learning: network distinctiveness predicts age of acquisition. *Cognitive Science*, 41, 120-140.
- Engelthaler, T., & Hills, T. T. (2019). Modeling early word learning through network graphs. In M. S. Vitevitch (Ed.), *Network Science in Cognitive Psychology* (pp. 166-183): Routledge.
- Estes, W. K. (1956). The problem of inference from curves based on group data. *Psychological Bulletin*, 53(2), 134.
- Estes, W. K. (1975). Some targets for mathematical psychology. *Journal of Mathematical Psychology*, 12(3), 263-282.
- Forster, K. I. (1978). Accessing the mental lexicon. In E. Walker (Ed.), *Explorations in the Biology of Language*. Montgomery, VT: Bradford.
- Freud, S. (1920). *A general introduction to psychoanalysis*: Createspace Independent Publishing Platform.
- Gershman, S. J. (2019). The generative adversarial brain. *Frontiers in Artificial Intelligence*, 2(18).
- Goñi, J., Martincorena, I., Corominas-Murtra, B., Arrondo, G., Ardanza-Trevijano, S., & Villoslada, P. (2010). Switcher-random-walks: A cognitive-inspired mechanism for network exploration. *International Journal of Bifurcation and Chaos*, 20(3), 913-922.
- He, L., Kenett, Y. N., Zhuang, K., Liu, C., Zeng, R., Yan, T., Huo, T., & Qiu, J. (2020). The relation between semantic memory structure, associative abilities, and verbal and figural creativity. *Thinking & Reasoning*, 1-26.
- Hills, T. T. (2006). Animal foraging and the evolution of goal-directed cognition. *Cognitive Science*, 30(1), 3-41.
- Hills, T. T. (2019). Neurocognitive free will. *Proceedings of the Royal Society B*, 286(1908), 20190510.
- Hills, T. T., Jones, M. N., & Todd, P. M. (2012). Optimal foraging in semantic memory. *Psychological Review*, 119(2), 431-440.
- Hills, T. T., Maouene, J., Riordan, B., & Smith, L. B. (2010). The associative structure of language: Contextual diversity in early word learning. *Journal of Memory and Language*, 63(3), 259-273.
- Hills, T. T., Maouene, M., Maouene, J., Sheya, A., & Smith, L. (2009a). Categorical structure among shared features in networks of early-learned nouns. *Cognition*, 112(3), 381-396.
- Hills, T. T., Maouene, M., Maouene, J., Sheya, A., & Smith, L. (2009b). Longitudinal analysis of early semantic networks: preferential attachment or preferential acquisition? *Psychological Science*, 20(6), 729-739.
- Hills, T. T., Mata, R., Wilke, A., & Samanez-Larkin, G. R. (2013). Mechanisms of age-related decline in memory search across the adult life span. *Developmental Psychology*, 49(12), 2396.
- Hills, T. T., & Pachur, T. (2012). Dynamic search and working memory in social recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 38(1), 218.
- Hills, T. T., & Segev, E. (2014). The news is American but our memories are... Chinese? *Journal of the Association for Information Science and Technology*, 65(9), 1810-1819.
- Hills, T. T., & Siew, C. S. Q. (2018). Filling gaps in early word learning. *Nature Human Behaviour*, 2(9), 622-623.
- Hintzman, D. L. (1974). Psychology and the cow's belly. *The Worm Runner's Digest*, 16, 84-85.

- Holtmaat, A., & Svoboda, K. (2009). Experience-dependent structural synaptic plasticity in the mammalian brain. *Nature Reviews Neuroscience*, 10(9), 647-658.
- Jagiello, R. D., & Hills, T. T. (2018). Bad news has wings: Dread risk mediates social amplification in risk communication. *Risk Analysis*, 38(10), 2193-2207.
- James, W. (1890). *The principles of psychology*. Cambridge, MA: Harvard University Press.
- Jamieson, R. K., Avery, J. E., Johns, B. T., & Jones, M. N. (2018). An instance theory of semantic memory. *Computational Brain & Behavior*, 1(2), 119-136.
- Jones, M., & Love, B. C. (2011). Bayesian fundamentalism or enlightenment? On the explanatory status and theoretical contributions of Bayesian models of cognition. *Behavioral and Brain Sciences*, 34(4), 169-231.
- Jones, M. N., Hills, T. T., & Todd, P. M. (2015). Hidden processes in structural representations: A reply to Abbott, Austerweil, and Griffiths (2015). *Psychological Review*, 122(3), 570-574.
- Kahneman, D. (2011). *Thinking, fast and slow*. Macmillan.
- Karras, T., Laine, S., & Aila, T. (2019). *A style-based generator architecture for generative adversarial networks*. Paper presented at the Proceedings of the IEEE conference on computer vision and pattern recognition.
- Karuza, E. A., Kahn, A. E., Thompson-Schill, S. L., & Bassett, D. S. (2017). Process reveals structure: How a network is traversed mediates expectations about its architecture. *Scientific Reports*, 7(1), 12733.
- Kenett, Y. N. (2018). Going the extra creative mile: The role of semantic distance in creativity – theory, research, and measurement. In R. E. Jung & O. Vartanian (Eds.), *The Cambridge Handbook of the Neuroscience of Creativity* (pp. 233-248). New York, NY: Cambridge 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., Beaty, R. E., Silvia, P. J., Anaki, D., & Faust, M. (2016). Structure and flexibility: Investigating the relation between the structure of the mental lexicon, fluid intelligence, and creative achievement. *Psychology of Aesthetics, Creativity, and the Arts*, 10(4), 377-388.
- Kenett, Y. N., & Faust, M. (2019a). Clinical cognitive networks: A graph theory approach. In M. S. Vitevitch (Ed.), *Network Science in Cognitive Science* (pp. 136-165). New York, NY: Routledge.
- Kenett, Y. N., & Faust, M. (2019b). A semantic network cartography of the creative mind. *Trends in Cognitive Sciences*, 23(4), 271-274.
- Kenett, Y. N., Kenett, D. Y., Ben-Jacob, E., & Faust, M. (2011). Global and local features of semantic networks: Evidence from the Hebrew mental lexicon. *PLoS ONE*, 6(8), e23912.
- Kenett, Y. N., Levi, E., Anaki, D., & Faust, M. (2017). The semantic distance task: Quantifying semantic distance with semantic network path length. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 43(9), 1470-1489.
- 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. (2021). Semantic memory: A review of methods, models, and current challenges. *Psychonomic Bulletin & Review*, 28(1), 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.

- Lambon Ralph, M. A., Jefferies, E., Patterson, K., & Rogers, T. T. (2017). The neural and computational bases of semantic cognition. *Nature Reviews Neuroscience*, 18(1), 42-55.
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436-444.
- Lerner, I., Bentin, S., & Shriki, O. (2012). Spreading activation in an attractor network with latching dynamics: Automatic semantic priming revisited. *Cognitive Science*, 36(8), 1339-1382.
- Lerner, I., Bentin, S., & Shriki, O. (2014). Integrating the automatic and the controlled: Strategies in semantic priming in an attractor network with latching dynamics. *Cognitive Science*, 38(8), 1562-1603.
- Lerner, I., & Shriki, O. (2014). Internally- and externally-driven network transitions as a basis for automatic and strategic processes in semantic priming: Theory and experimental validation. *Frontiers in Psychology*, 5, 314.
- Lurie, D. J., Kessler, D., Bassett, D. S., Betzel, R. F., Breakspear, M., Kheilholz, S., Kucyi, A., Liégeois, R., Lindquist, M. A., McIntosh, A. R., Poldrack, R. A., Shine, J. M., Thompson, W. H., Bielczyk, N. Z., Douw, L., Kraft, D., Miller, R. L., Muthuraman, M., Pasquini, L., Razi, A., Vidaurre, D., Xie, H., & Calhoun, V. D. (2019). Questions and controversies in the study of time-varying functional connectivity in resting fMRI. *Network Neuroscience*, 4(1), 30-69.
- McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: I. An account of basic findings. *Psychological Review*, 88(5), 375-407.
- McKoon, G., & Ratcliff, R. (1992). Spreading activation versus compound cue accounts of priming: Mediated priming revisited. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18(6), 1155.
- Morais, A. S., Olsson, H., & Schooler, L. J. (2013). Mapping the structure of semantic memory. *Cognitive Science*, 37(1), 125-145.
- Moussaïd, M., Brighton, H., & Gaissmaier, W. (2015). The amplification of risk in experimental diffusion chains. *Proceedings of the National Academy of Sciences*, 112(18), 5631-5636.
- Musz, E., & Thompson-Schill, S. L. (2015). Semantic variability predicts neural variability of object concepts. *Neuropsychologia*, 76, 41-51.
- Myung, I. J., Kim, C., & Pitt, M. A. (2000). Toward an explanation of the power law artifact: Insights from response surface analysis. *Memory & Cognition*, 28(5), 832-840.
- 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.
- Park, D. C., Lautenschlager, G., Hedden, T., Davidson, N. S., Smith, A. D., & Smith, P. K. (2002). Models of visuospatial and verbal memory across the adult life span. *Psychology and Aging*, 17(2), 299.
- Posner, M. I., & Konick, A. F. (1966). On the role of interference in short-term retention. *Journal of Experimental Psychology*, 72(2), 221-231.
- Roediger, H. L. (1980). Memory metaphors in cognitive psychology. *Memory & Cognition*, 8(3), 231-246.
- Rogers, T. T., Patterson, K., Jefferies, E., & Lambone-Ralph, M. A. (2015). Disorders of representation and control in semantic cognition: Effects of familiarity, typicality, and specificity. *Neuropsychologia*, 76, 220-239.
- Rosen, V. M., & Engle, R. W. (1997). The role of working memory capacity in retrieval. *Journal of Experimental Psychology: General*, 126(3), 211-227.
- Salthouse, T. A., Atkinson, T. M., & Berish, D. E. (2003). Executive functioning as a potential mediator of age-related cognitive decline in normal adults. *Journal of Experimental Psychology: General*, 132(4), 566.

- Shiffrin, R. M., Lee, M. D., Kim, W., & Wagenmakers, E. J. (2008). A survey of model evaluation approaches with a tutorial on hierarchical Bayesian methods. *Cognitive Science*, 32(8), 1248-1284.
- Siew, C. S. Q. (2018). The orthographic similarity structure of English words: Insights from network science. *Applied Network Science*, 3(1), 13.
- Siew, C. S. Q., Engelthaler, T., & Hills, T. T. (under review). Nymph piss and gravy orgies: Local and global contrast effects in relational humor.
- Siew, C. S. Q., 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, 24.
- Silvia, P. J. (2015). Intelligence and creativity are pretty similar after all. *Educational Psychology Review*, 27(4), 1-8.
- Sizemore, A. E., Karuza, E. A., Giusti, C., & Bassett, D. S. (2017). Knowledge gaps in the early growth of semantic networks. *arXiv preprint*.
- Stella, M. (2019). Modelling early word acquisition through multiplex lexical networks and machine learning. *Big Data and Cognitive Computing*, 3(1), 10.
- Stella, M. (2020). Multiplex networks quantify robustness of the mental lexicon to catastrophic concept failures, aphasic degradation and ageing. *Physica A: Statistical Mechanics and its Applications*, 554, 124382.
- Stella, M., Beckage, N. M., & Brede, M. (2017). Multiplex lexical networks reveal patterns in early word acquisition in children. *Scientific Reports*, 7, 46730.
- Stella, M., Beckage, N. M., Brede, M., & De Domenico, M. (2018). Multiplex model of mental lexicon reveals explosive learning in humans. *Scientific Reports*, 8(1), 2259.
- Stella, M., & Brede, M. (2016). Mental lexicon growth modelling reveals the multiplexity of the English language. In H. Cherifi & B. Gonçalves & R. Menezes & R. Sinatra (Eds.), *Complex Networks VII: Proceedings of the 7th Workshop on Complex Networks CompleNet 2016* (pp. 267-279). Cham: Springer International Publishing.
- Stella, M., & Kenett, Y. N. (2019). Viability in multiplex lexical networks and machine learning characterizes human creativity. *Big Data and Cognitive Computing*, 3(3), 45.
- Steyvers, M., & Tenenbaum, J. B. (2005). The large scale structure of semantic networks: Statistical analysis and a model of semantic growth. *Cognitive Science*, 29(1), 41-78.
- Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: Statistics, structure, and abstraction. *Science*, 331(6022), 1279-1285.
- Tolman, E., & Tolman, E. C. (2004). *The Psychological Review Cognitive Maps in Rats and Men* *.
- Turvey, M. T., & Moreno, M. A. (2006). Physical metaphors for the mental lexicon. *The Mental Lexicon*, 1(1), 7-33.
- Vitevitch, M. S. (2008). What can graph theory tell us about word learning and lexical retrieval? *Journal of Speech Language and Hearing Research*, 51(2), 408-422.
- Vitevitch, M. S., & Castro, N. (2015). Using network science in the language sciences and clinic. *International Journal of Speech-Language Pathology*, 17(1), 13-25.
- Vitevitch, M. S., Chan, K. Y., & Goldstein, R. (2014). Insights into failed lexical retrieval from network science. *Cognitive Psychology*, 68, 1-32.
- Volle, E. (2018). Associative and controlled cognition in divergent thinking: Theoretical, experimental, neuroimaging evidence, and new directions. In R. E. Jung & O. Vartanian (Eds.), *The Cambridge handbook of the neuroscience of creativity* (pp. 333-362). New York, NY: Cambridge University Press.
- Wulff, D. U., De Deyne, S., Jones, M. N., Mata, R., & Consortium, T. A. L. (2019). New perspectives on the aging lexicon. *Trends in Cognitive Sciences*, 23(8), 686-698.
- Wulff, D. U., Hills, T., & Hertwig, R. (2020). Memory is one representation not many: Evidence against wormholes in memory. *arXiv*.

- Wulff, D. U., Hills, T., & Mata, R. (2018). Structural differences in the semantic networks of younger and older adults. *PsyArXiv*.
- Yee, E., & Thompson-Schill, S. L. (2016). Putting concepts into context. *Psychonomic Bulletin & Review*, 23(4), 1015-1027.
- Zemla, J. C., & Austerweil, J. L. (2018). Estimating semantic networks of groups and individuals from fluency data. *Computational Brain and Behavior*, 1(1), 36-58.
- Zemla, J. C., Kenett, Y. N., Jun, K.-S., & Austerweil, J. L. (2016). U-INVITE: Estimating individual semantic networks from fluency data. In A. Papafragou & D. Grodner & D. Mirman & J. C. Trueswell (Eds.), *Proceedings of the 38th Annual Meeting of the Cognitive Science Society* (pp. 1907-1912). Austin, TX: Cognitive Science Society.
- Zortea, M., Menegola, B., Villavicencio, A., & Salles, J. F. d. (2014). Graph analysis of semantic word association among children, adults, and the elderly. *Psicologia: Reflexão e Crítica*, 27, 90-99.