20

On Cognition

*Network Representation Capacity: How Social
Relationships are Represented in the Human Mind*

Hui Sun, Matthew E. Brashears, and Edward B. Smith

Collecting network data with an egocentric approach – by eliciting respondent's (i.e., ego's) self-report of her direct contacts (i.e., alters) – is a popular tactic in research on communities (e.g., Fischer et al. 1977), health/distress (e.g., The National Longitudinal Study of Adolescent to Adult Health), and general social behavior (e.g., General Social Survey). However, accuracies of self-reported networks are heavily debated. On the one hand, a plethora of evidence has documented the extent of inaccuracy in egos' reports of their own network ties (e.g., Bernard, Killworth, and Sailer 1982), especially when the relations being assessed are more nuanced than "knowing" (e.g., Krackhardt 1996; Kumbasar, Rommey, and Batchelder 1994; Krackhardt and Kilduff 1999). For example, Krackhardt (2014) observed that for twelve out of the twenty-one managers from a small entrepreneurial firm, their beliefs about tie existence were wrong more than half of the time. From a methodological perspective, this leads to substantial doubt about the validity of self-reported egocentric networks; errors in individual reports can have a material impact on the structural inference we draw from data, such as the centralization of the network and the ego's structural position within it (Borgatti, Carley, and Krackhardt 2006; Krackhardt 2014).

On the other hand, accumulating evidence suggests that despite individuals' inaccuracy in perceiving local network structures (i.e., dyadic ties), people tend to have a fairly good idea of the global network structure (i.e., larger network structure of all relationships; Burt and Bittner 1981; Casciaro, Carley, and Krackhardt 1999). For instance, although centrality measures (especially eigenvector centrality) derived from self-reported egocentric networks have been criticized as biased, individuals often demonstrate superb ability to track popularity in real-world social networks (Zerubavel et al. 2015).[1] In a study

---

[1] Developmental psychology research suggests that such abilities may have developed as early as preschool age (Vaughn and Waters 1981; Lansu, Cillessen, and Karremans 2014).

investigating how affect influences accuracy in network perception, Casciaro and her colleagues (1999) found that while positive affect is positively correlated with global accuracy, its association with local accuracy appears to go in the opposite direction. The puzzling dissociation between local inaccuracy and global accuracy demands a serious theoretical inquiry: how are social networks represented in the mind of an individual? If individuals are indeed poor at perceiving the structure, particularly the local structure, of networks in which they are embedded, how is the mobilization of social capital even possible (e.g., Granovetter 1973; Lin 1999; Burt 2009)? Alternatively, if individuals are knowledgeable about the global network structure, what can account for high error rates in their perceptions[2] of local network structures?

Research in cognitive science may provide a clue. Converging evidence suggests that human representation of social networks varies in levels of abstraction (Rumelhart and Ortony 1978; Kemp and Tenenbaum 2008). It is very likely that our perceptions of local network features are encoded and retrieved through different relational schemata – mental structures of how social agents are expected to be structurally connected – in comparison to perceptions of global network structures. In other words, the apparent puzzle between local inaccuracy and global accuracy could indicate separate but related schemata are employed simultaneously at different levels of abstraction (i.e., global versus local). Within each schema complex network information with varying granularity (i.e., more or less detailed) is encoded and stored. In a way, our view is congruent with emerging evidence in egocentric network research that discrepancies in elicited networks may be more substantive than previously conceived, reflecting cognitive underpinnings in addition to more prosaic methodological concerns (Pescosolido and Wright 2004).

Our theory of a multi-schemata representation of social networks deviates from those that assume either a local or global representation of social networks. In our view, the tension between local inaccuracy and global accuracy is a "symptom" of the scholar's atomistic assumption that the ego's knowledge of social relationships is a composite of her perceptions of individual ties. Because social networks are analytically constructed from dyadic relationships, the discovery process sometimes gives us the impression that perceptions of social networks should also take a bottom-up approach, with perceptions of higher structures rooted in knowledge of lower-level structures. Nonetheless, even though tie perception is a logical antecedent to network perception, it does not guarantee that a summary statistic of accuracy aggregated from tie perception, weighted or not, can accurately reflect the psychological reality of one's knowledge of social networks. Through

---

[2] Perception here refers to a general understanding of a target elicited from subjective reports. The scope of perception in this chapter is akin to social psychologists' use of perception, rather than the strict definition of perception typically found in cognitive psychology literature.

experience, ~~human~~ individuals may develop an ecologically reasonable perception of the higher-level structure of a network that is schematically distinct from any lower-level representation, and therefore only loosely correlates with it.

Meanwhile, we also differ from perspectives that assume a holistic primacy, meaning that global representation should override local ones. Empirical nuances notwithstanding (Alba and Hasher 1983; Wagemans et al. 2012), theoretical interpretations of Gestalt psychology often imply that human cognition of objects, concepts and social relationships operates in a holistic fashion (Brewer and Nakamura 1984). Rather than integrating featural components to form a representation of the whole (e.g., Treisman 1986), individuals are believed to actively construct a representation of the whole and then infer properties of components from it (Kimchi 2003). In this view, local features of a network are derived from this global abstraction and therefore are in some sense subservient to it.

The tension between the two approaches is likely to be exaggerated in scholarly debate, but it reflects a theoretical distinction between two schools of thought on human mental representations of knowledge. The classical view (e.g., Newell and Simon 1972; Fodor and Pylyshyn 1988), partly derived from research on languages, holds that knowledge is mentally represented as rule-based, symbolic structures. One key characteristic of this type of structure is that representations are compositional, such that a representation of the whole is determined by the contents of its constituents and their structural configuration. The connectionist view, on the other hand, sees representations as networks of interconnected nodes (Rumelhart and Mcclelland 1986; Smolensky 1988). Under this framework, representations are the result of parallel processing and therefore they are intrinsically holistic. No single node possesses any meaning on its own; meanings are derived from its connection with other constituents. There is as yet no compelling evidence that one school of thought is always superior than the other. Rather, they are best regarded as two ideal types for describing, prescribing and predicting cognitive behavior, meaning that they may both be correct to varying extents for particular types of processing. Preference for one or the other, however, is still a substantial decision as it constrains the questions we will ask, the theories we can use, and the computational models we may develop.

From our perspective, an individual's accuracy in reporting network characteristics reflects her mental representation of the social network. These representations consist of collections of schemata adopted by the individual to encode and store complex network information. These schemata operate at different levels of granularity (i.e., more or less detailed) and abstraction (i.e., global versus local) and can be active simultaneously. They are used for dealing with different types of concrete problems and the inconsistency between them accounts for varying accuracy in the individual's perception of global and local network properties.

NETWORK REPRESENTATIONS AND RELATIONAL SCHEMATA

In broad strokes, social networks and communication networks are highly patterned (Albert and Barabási 2002; Ebel, Mielsch, and Bornholdt 2002; Newman and Park 2003). Local features such as triadic closure (i.e., the propensity of pairs of actors to be connected when they share a common neighbor), and assortative mixing (i.e., tendency for connected nodes to have similar numbers of ties) are typical of social networks due to the prevalence of community structures in such networks (Newman and Park 2003). From the local interactions of a large number of individuals, global regularities such as small-world properties (Watts and Strogatz 1998) and scale-free or log-normal distributions of the degree (Barabási and Albert 1999; Broido and Clauset 2019) frequently emerge.

Therefore, it is not surprising that our mind adapts to these regularities in our social relational environment by forming mental structures that organize information and by developing mental processes that act on these structures (Chase and Simon 1973; Alba and Hasher 1983; Hintzman 1986; Markman 1999). Typically referred to as schemata, these mental structures are frameworks of systematic knowledge that individuals acquire over time (Bartlett, 1932) to guide the selection, abstraction, interpretation, integration, and consolidation of new information (Alba and Hasher 1983).[3] They enrich the manifest stimulus by enabling individuals to go beyond the currently available information, filling in the blanks in what is observed with information derived from the schema (Fiedler 1982).

Here, following Kelley's (1972) approach to defining causal schema, we define relational schemata as conceptions of the manner in which social contacts are structurally related. These relational schemata collectively form an individual's generic knowledge of the structure in her social environment, and changes to the social environment (e.g., when the individual learns about a new person) may trigger updates in this knowledge.

From a functional perspective, schemata have two benefits. First, the use of schemata reduces the burden of cognitive processing. As Brashears (2013) showed in his experiments, humans often adaptively use schemata to compress complex network information so that they can store large networks with finite cognitive resources (see also Brashears and Brashears 2020). Second, the existence of schemata makes generalization possible. The ability to generalize has always been seen as a crucial component of human intelligence.

---

[3] Scholars' definitions of schemata and their corresponding ontological assumptions vary widely (Brewer and Nakamura 1984). Bartlett (1932: 201) defined a schema as "an active organization of past reactions, or of past experiences, which must always be supposed to be operating in any well-adapted organic response". Minsky (1975) and Rumelhart and Ortony (1978: 101) took a more instrumentalist position and defined them as data structures. Neisser (1976: 54) took a realist position and defined a schema as "a part of the nervous system". In social network traditions, schemas are typically referred to as mental models.

While artificial intelligence agents require hundreds of thousands of data points to learn a new concept and make correct judgments on previously unseen stimuli, a human child can accomplish the same task with only a few examples. Although exactly how humans have such special aptitudes for forming schemata of the environment is a topic of active research (Holyoak 2008; Kemp and Tenenbaum 2008; Schapiro and Turk-Browne 2015), there is no doubt that schemata are critical for making inferences and setting expectations in uncertain or novel situations, as well as that humans are unusually good at doing so (Holland et al. 1989).

As pointed out by Rumelhart and Ortony (1978), schemata are often hierarchical. Relational schemata are no exception. The system of relational schemata that an individual possesses include both schemata of local dyadic relationships (Karuza, Thompson-Schill, and Bassett 2016) as well as schemata of global relational structures (Parkinson, Kleinbaum, and Wheatley 2017; Tompson et al. 2019). For instance, research participants exhibit a tendency to infer triadic closure when there is none, illustrating the role of a local relational schema in leading participants to go beyond the given information and "fill in the blanks" in a structure (Freeman 1992). Empirical evidence of global relational schemata is comparatively less known to network scholars. A few notable examples include Bond, Jones, and Weintraub's (1985) proposition that social contacts are mentally organized into social-group categories; Bond and Brockett's (1987) social context-personality index theory that social contacts are first organized according to their encoding contexts (i.e., when and where each alter was encountered), and then further partitioned into different personality sub-clusters; and Fiske's (1995) and Brewer's (1995b) position that mental organization is framed by interaction and affiliation patterns, thereby more closely resembling modern social network theory.

Recent effort to connect social networks with cognitive neuroscience has opened up new opportunities to revisit these old questions. For example, Parkinson and colleagues (2017) paired fMRI scans of a group of MBA participants with measurement of their social networks, and found that participants spontaneously encoded other cohort members' social network positions (e.g., degrees of separation from themselves, eigenvector centrality and brokerage) when seeing their faces. Despite the absence of a network visualization, participants in the study automatically mapped alters onto a network structure in their mind, suggesting that they have adopted a global schema of social relations that is very close to a network structure. However, in some cases, participants may also adopt a global schema that deviates substantially from a network structure. In one study, Tavares and colleagues (2015) asked participants to play lead characters in a role-playing game while scanning their brains. They found that participants' global schema of social relations was best characterized as a two-dimensional social space framed by power and affiliation, instead of a network

characterized by social interactions. In short, while human brains seem predisposed to spontaneously encode social relationships into a structural format, it does not appear to be inevitable that this structural format resemble what we typically consider as a network.

In the context of network representation – one's mental representation of social networks – schemata may also come at a cost. Although schemata enable individuals to represent large amounts of social relational information with limited mental space, they run the risk of deviating from the network that they seek to represent. Some deviations can be attributed to individuals' lack of experience with some forms of network structures and therefore these structures have not been fully incorporated into schemata (Winkler-Rhoades et al. 2010). For example, prior experience with local network features is found to influence how one represents a new network structure. When participants were asked to learn a new network structure that they had less experience with, their learning performance was impaired (Janicik and Larrick 2005).

Some deviations are due to contextual configurations of the task that can influence an individual's network representation by shifting the schema she adopts. For instance, it was found that adding kinship labels to ties facilitated learning when the network structure to be learned corresponded with common expectations of a kinship network (Brashears 2013). Likewise, individuals recall imbalanced affect in networks more accurately when a source of constraint is able to account for the persistence of this imbalance (Brashears and Brashears 2016). In short, when an appropriate schema is activated, the gap between the network structure and one's network representation is small, but when an inappropriate schema is activated, this gap can be quite large.

PARADIGMS FOR QUANTIFYING INDIVIDUAL DIFFERENCES IN NETWORK REPRESENTATIONS

Given that individuals often have varying experience with different types of network structures, and that they may have heterogeneous preferences for schemata in various contexts, the adequacy and appropriateness of the schemata they adopt to represent their social network in a given context may also vary. However, it is an empirical challenge to discern the exact relational schema recruited by participants using the egocentric networks they report because the same network data could have been elicited via several possible schemata. Participants may represent social relationships as multidimensional social networks (John Scott 1988; Fiori, Smith, and Antonucci 2007), as groups (Gershman, Pouncy, and Gweon 2017) and social circles (Hill and Dunbar 2003; Zhou et al. 2005), as a lower-dimensional vector space of alters (Tang et al. 2015), or even in a form that has not yet been considered. Similarly, they may represent psychological distance among individuals as geodesic distance (Parkinson et al. 2017), as Euclidean distance over a multi-dimensional feature

space (Tavares et al. 2015), as the extent of trait similarities (Carlston and Skowronski 1994; Uleman, Saribay, and Gonzalez 2007), or via some entirely different approach. In short, inferring which schema was used from a given measured network is challenging because we only obtain the end state, without obvious information on the preceding cognitive processes.

Despite the difficulty in discerning the exact relational schema individuals use to organize social information in a given situation for a specific goal, we are not without options. One option is to quantify the gap between an individual's social network and her mental representation of it, the result of which reflects what we term their *the network representation capacity*. However, higher network representation capacity may not directly lead to better mobilization of social network resources; across many domains of life, knowledge of the right thing to do is often distinct from the propensity or ability to actually do it. It is likely that a coarser representation, or lower correspondence between one's network representation and her social networks, might generate better utility in some scenarios.

QUANTITATIVE PARADIGMS FOR NETWORK REPRESENTATION

Four survey or experimental paradigms can be used to quantify individuals' network representation capacities: the error paradigm, the free-recall paradigm, the structural learning paradigm and the statistical learning paradigm. Both the error paradigm and the free-recall paradigm rely on memory principles: the error paradigm, as its name reveals, makes use of observed errors during recall while the free-recall paradigm takes advantage of recall sequences to infer the gap between the schematic organization and the social network organization of social contacts. In contrast, the structural learning and the statistical learning paradigms rely on behavioral indicators when learning new network structures. The rationale is that if one uses a specific set of schemata to organize her current egocentric networks, she is more likely to apply the same set of schemata to a new network under similar context.

Error Paradigm

Errors often reflect the schemata that we depend on when encoding, representing and retrieving social relational information. An incorrect perception of tie existence can be attributed to selective encoding (Minsky 1975; Schank and Abelson 1977), abstraction, interpretation, or integration of information (Sulin and Dooling 1974; Minsky 1975; Schank and Abelson 1977; Wood 1978). For example, we may selectively ignore a waiter's brief nod to his friend in a restaurant, often without being consciously aware of having done so, because of the schematic assumption of a service relationship, and accordingly misclassify the relationship as acquaintances or even strangers. Likewise, we may mistakenly draw the inference that

friend A and friend B know each other after finding both of them posting pictures of the same concert, even though in reality they have never come across each other. Freeman (1992) notably described this process as "filling in the blanks," while Harris and Monaco (1978) referred to it as inference.

The error paradigm takes advantage of the mental processes that occur during information processing and derives inferences of network representation capacity through statistical analysis of recall errors. For example, measures of commission errors (i.e., inferring existence of a feature when there is none) and omission errors (i.e., reporting nonexistence of a feature when there is one) can capture an individual's network representation capacity for dyadic relations within a network (Krackhardt 2014). These two measures can also be extended to bundles of ties. For instance, we can develop similar measures of individuals' tendency to infer transitivity and reciprocity when no other information about the network is provided (De Soto, Henley, and London 1968; Crockett 1982; Krackhardt and Kilduff 1999). In short, network representation capacity can be operationalized as performance of a binary classifier under the error paradigm. Naturally, comparisons can be drawn based on measures of discriminant ability such as a ROC ~~curve~~ *curves* (Yonelinas 1994, 2002).

Although we are usually more interested in representation of network structures than representation of nodes, the error paradigm may also be extended to misattribution of node characteristics. For instance, it is common for individuals to infer similarities in taste based on social connections (Berger and Heath 2008), or similarities in some irrelevant dimensions (Todorov and Uleman 2003). Therefore, errors may also arise in recalling node characteristics (e.g., trait, preferences, attitudes) because people often overestimate the extent of similarities among connected others in our network (Prentice 1990).

One limitation of the error paradigm comes from a lack of "ground truth." When two individuals are compared based on their network representation capacity for a specific network feature, an equivalence needs to be established such that the networks they are representing are of comparable difficulty. One remedy is to normalize the resulting measure based on the size of the network as well as the total number of possible ties or triads. An alternative solution is to incorporate experimental methods and directly manipulate the network to be represented. In that case, the experimenter may design a network structure to be recalled so that the accuracy of a representation can be definitively measured.

In general, errors are direct indicators of one's network representation capacity. However, it is mostly applicable to representations of local network features. To draw comparisons across individuals, we first specify a particular structure to be compared upon (e.g., triadic closure), and then infer network representation capacity using a measure of discriminant ability based on error patterns.

Free-Recall Paradigm

The free-recall paradigm is similar to the error paradigm because both of them take advantage of mental processes, or more specifically memory processes, to infer the gap between social networks and individuals' mental representations of them. However, instead of focusing on representation of local network features, the free-recall paradigm is most powerful in capturing representations of global network structures.

The paradigm explicitly makes use of the fact that name generators (i.e., tasks that ask egos to list the names of their contacts) and name interpreters (i.e., tasks that collect information about those alters) are, by their nature, memory tasks. When respondents are cued about their social contacts, they activate memories of their social relationships. The resulting egocentric network is the output of the ego's schematic encoding of her social relationships into network representations and schematic retrieval from such representations. In other words, the seemingly inaccurate perceptions of ties are also patterned, organized reflections of the contents of individual memory (Bernard et al. 1982; Freeman, Romney, and Freeman 1987; Freeman 1992). Analogous to psychologists' response to economists' assumption of human rationality, recall errors in social networks are meaningful despite, and often because of, significant deviations from actual interaction history. In the end, it is retrieval from that representation – accurate or inaccurate – that becomes the basis for subsequent action.

A typical free-recall task (also referred to as a free-listing task, an unconstrained recall task or a fluency task) either asks participants to freely generate a specific number of items that fall under given criteria, or to generate any number of items that fall under given criteria within a pre-specified period of time (e.g., sixty seconds). Using free recall to infer an underlying network representation relies on one crucial memory principle: items that are retrieved consecutively are often closer in mental space, because people have a natural tendency to organize information into meaningful clusters. For example, in one of their experiments, Hamilton, Leirer, and Katz (1979) presented participants with behavioral descriptions from four categories in random order. However, during recall sessions, participants naturally grouped the behavior into four clusters without being explicitly told to do so. Similar tendencies are observed in experiments that ask participants to generate names of United States cities (Bousfield and Sedgewick 1944), food (Gruenewald and Lockhead 1980), occupations (Graesser and Mandler 1978) and people (Bond and Brockett 1987).

Such robust memory tendencies present a unique opportunity to uncover how well one's network representation maps onto the reported egocentric network.[4] Ideally, if a person's network representation maps perfectly to the

---

[4] Although it is also viable to use more "objective" measure of networks, such as those constructed from interaction data or aggregation of multiple people's self-reported network relations (e.g., Krackhardt 1987), we think it is preferable to use self-reported egocentric network as the comparison benchmark for two reasons. First, it isolates the effect due to deviations in knowledge

network she reports, we would anticipate that she will first deplete names that belong to one network cluster, before moving to another cluster (Roenker, Thompson, and Brown 1971). In other words, in a perfect match scenario, the activation path should be as short as possible. However, if the person generates in succession two names that have a moderate to large geodesic distance between them, it implies that these two contacts are close in the ego's mental space despite being distant in the reported network space. This could be attributed to "imperfect" network representation as the ego may have relied on some other dimensions to represent two people in the mental space in addition to their relational closeness (equivalent to inserting of ties, e.g., Brewer 1995a, 1995b, 1997; Brewer et al. 2005). Alternatively, the ego may have even adopted a different global relational schema, such as a hierarchical tree structure or a categorical structure instead of a network structure to represent social relationships.

Recently, two of this chapter's authors developed a measure – the jump ratio – based on this idea (Sun and Smith 2018). They first operationalized the "jump" as the sum of a sequence of distance measures, where each segment in the sequence is the geodesic distance between two consecutively recalled nodes. The intuition is that searching for contacts in one's mind is analogous to designing a travel route, the more the ego switches between clusters that are far apart, the more detours they take, and the longer the travel will be. To make the measure comparable across egos with networks of different size and connectivity, they then normalized jump by subtracting the shortest possible activation path that the ego could potentially traverse on her network and then divided the resulting value by the range between the longest possible activation path and the shortest possible activation path. Analogously, it is unfair to say that a global traveler is less efficient in designing her traveling route than a United States traveler simply because the former travels farther. To make these two travelers comparable, their traveling distance needs to be contextualized to the map that they traverse – how would they compare to a hypothetical traveler who travels on the same map but always picks the longest possible route? The resulting normalized measure is the "*jump ratio*." Formally, in an egocentric social network with n network contacts, V_i denotes the i th recalled alter, $d(V_i, V_{i+1})$ represents the geodesic distance between alter V_i and V_{i+1} minus 1,[5] and P refers to the set of all possible ways of recalling the

---

structure from incorrect perception of ties. If the ego already perceives two people as disconnected, but still recalls them in succession, it implies that the ego relies on a knowledge representation that may not be network based. Second, it has been found that people are more susceptible to commission errors (i.e., inserting ties when there are none) than omission errors (Krackhardt 2014). Therefore, recall path lengths calculated based on self-reported egocentric networks tend to underestimate the gap. In other words, it is a conservative test.

[5] When V_i and V_{i+1} are directly connected, the geodesic distance between them equals one. Subtracting geodesic distance by one ensures that in these cases, the mental distance traveled equals zero.

same *n* network contacts.[6] The jump ratio a person establishes when cognitively navigating their social network is thus akin to the total (social) distance traveled in the process of search normalized by the range of possible distance,[7] or:

$$JumpRatio = \frac{\sum_{i=1}^{n-1} d(V_i, V_{i+1}) - \min_{p \in P} \sum_{i=1}^{n-1} d(V_{p_i}, V_{p_i+1})}{\max_{p \in P} \sum_{i=1}^{n-1} d(V_{p_i}, V_{p_i+1}) - \min_{p \in P} \sum_{i=1}^{n-1} d(V_{p_i}, V_{p_i+1})} \quad \text{(A)}$$

In contrast to the other paradigms, the free-recall paradigm[8] does not require researchers to turn on the spotlight and pre-specify the local or global network structure they are interested in. This is a blessing as well as a curse. Measures such as the jump ratio make direct comparison of individuals' general network representation capacity possible. However, since the measure incorporates every aspect of the network, it also makes adjusting for measurement errors difficult. For example, it is likely that some measurement errors are more probable than others and such errors will impact the resulting measure that we obtain.

By calculating the difference between one's actual activation path and a hypothetical activation path that will be expected if one's network representation is perfectly in congruence with the reported network structure, we can determine the strength of an individual's network representation capacity.

Structural Learning Paradigm

The structural learning paradigm is an experimental design wherein participants are asked to learn the underlying structure of a group of hypothetical actors (e.g., De Soto 1960). In these experiments, participants are usually presented with statements about the relationship between two people (e.g., Alice and Bob are friends) one at a time and asked to indicate whether the statement is true or false. Initial responses are pure guesses since participants have no prior knowledge about the relationships among those hypothetical actors. However, as participants get more feedback about their guesses, they begin to develop some network representation of the actors' social connections.

---

[6] For simplicity, distances between isolates and other nodes (i.e., alters who are not connected to any other alters in the network) are treated as one plus diameter (i.e., largest geodesic distance of the network).

[7] Keen readers may recognize that the denominator in the function is a variant of the famous traveling salesman problem. Although algorithmic solutions to such NP-hard problems are impossible, computational solutions are robust and readily available given their importance in operations research.

[8] Notably, the free-recall paradigm is not limited to capturing the representation capacity of networks. With multidimensional scaling and more sophisticated computational models, scholars have used it to explore human representation of categories and concepts among others (Medin et al. 1997; Griffiths, Steyvers, and Tenenbaum 2007).

This paradigm has its origin in research on associative learning. It was found that when a series of stimuli always appear in a pre-determined temporal or spatial order unknown to the participants, the associational relations will very soon be acquired by the participants in their behavioral responses (Paivio 1969). In fact, neuroscience has provided evidence that learned associations can be encoded through neural selectivity: with repeated exposure to structured stimulus, neurons, especially those in the medial temporal lobe, begin to respond similarly to these stimuli (Messinger et al. 2001; Osada et al. 2008). Association learning is especially crucial in the acquisition of syntax (e.g., Saffran & Wilson 2003), phonemes (Onnis, Waterfall, and Edelman 2008) and phonetics (e.g., Chambers, Onishi, and Fisher 2003).

Typically, those who learn the new structure faster (i.e., get all statements correct within fewer trials) are seen as having higher network representation capacity (e.g., Janicik and Larrick 2005). The rationale is that effective learning depends on proper use of schemata. Thus, faster learning indicates more expertise and better network representation of the structure, especially when the information to be learned is meaningful (Saariluoma 1989). In addition to number of trials, better network representation can also be detected through earlier recognition of regularities (e.g., Fiser and Aslin 2001) or faster responses (e.g., Kim et al. 2009) and/or more attention (e.g., Chun and Jiang 1999; Kidd, Piantadosi, and Aslin 2012) towards anticipated stimuli.

However, a comprehensive measure of network representation capacity requires combining learning performance over a portfolio of network structures. Because individuals' relational schemata are acquired through experience (Gobet and Simon 1996) and therefore their general capacity to represent networks will vary based on their prior experience with the specific network structure to be learned. When the network structure to be learned is similar to one's own network, we should expect better learning performance. This is exactly what Janicik and Larrick (2005) found in their experiment; participants whose own networks had open triads (i.e., when A is connected to B and B is connected to C, but A is not connected to C) were better able to learn networks that also had open triads.

Notably, the structural learning paradigm can be combined with the error approach to make comparisons more rigorous. In such scenarios, rather than using the number of trials as the measure of capacity, scholars turn towards patterns of errors in participants' recall of dyadic relationships between the hypothetical actors that they have previously seen. As an example, in several recent studies (Brashears 2013; Brashears and Quintane 2015), researchers found that the prevalence of "erroneously closed" ties varies depending on the schema activated during encoding and recall. The combination opens up new opportunities to investigate contextual and situational antecedents of network knowledge representation.

Statistical Learning Paradigm

Similar to the associative learning that underlies the structural learning paradigm, statistical learning[9] as a mental process, also reflects the human ability to "extract regularities from the environment" (Schapiro and Turk-Browne 2015, p.501). Statistical learning takes place throughout one's life span, from infancy (e.g., Kirkham, Slemmer, and Johnson 2002) to old age (e.g., Schapiro et al. 2014). It operates over multiple modalities and is widely recognized as a pervasive element of cognition (Saffran, Aslin, and Newport 1996; Fiser and Aslin 2001; Conway and Christiansen 2005; Brady and Oliva 2008; Gebhart, Aslin, and Newport 2009). In essence, it is the crucial cognitive function that supports the formation of representations of our physical and social environment, which naturally includes representation of social networks.

The statistical learning paradigm, relying on statistical learning as its cognitive underpinning, can be seen as an extension of the structural learning paradigm. Both of them rely on behavioral and/or neural indicators observed in a controlled, artificial learning task to infer relational schemata recruited by participants to facilitate learning. And both of them can rely on errors or learning time as a behavioral indicator of the prominence of a specific schema.

However, there are two key differences between the two paradigms. First, in contrast to the structural learning paradigm where learning performance is measured explicitly through accuracy in judgment of learned relations, in the statistical learning paradigm, learning is assumed to be automatic and implicit. It is assumed that individuals who are embedded in an environment gradually behave and believe in accordance with its underlying regularities even without conscious awareness. Second, the two paradigms differ in the structure of the world they construct for the participants. While both are learning tasks, the structural learning paradigm deterministically presents relationships to the participants while the statistical learning paradigm requires participants to infer relationships from probabilistic patterns among presented items.

Specifically, a structural learning task relies on repeated presentations of deterministic information on social relationships (e.g., "Alice and Bob are friends"), while a statistical learning task takes a more probabilistic approach either by presenting instances of interactions or by presenting a series of stimuli generated from a probabilistic transition matrix to suggest the existence of a relationship pattern. The general idea is that pairs of actors who interact more often are, ceteris paribus, more likely to have a closer relationship. And similarly, groups of actors who appear together more often are, ceteris paribus, more likely to be related to some global network structure, such as a community. In one of the recent studies that adopted the statistical learning

---

[9] The term "statistical learning" also has other uses in computer science (see Vapnik 1999), where it refers to theoretical and algorithmic analysis of function estimation. Many of the machine-learning algorithms that we are familiar with (e.g., support vector machine) originate in this definition of statistical learning.

paradigm, Tompson and colleagues (2019) were interested in participants' mental representation of community structure in social and non-social contexts. Under the social context condition, they presented participants with a series of fractal images one at a time and framed these images as avatars of people from an online social media platform. The task for participants was to decide whether each image was rotated or not. Unbeknownst to the participants, the sequence of images was generated by a random walk through a network structure pre-specified by the researchers, and each fractal image represented a node in the network. This random walk process ensured that after one fractal image was presented, its connected images were equally likely to be presented immediately afterwards.

To measure participants' capacity to learn the underlying network structure, Tompson and colleagues relied on two behavioral indicators. One is an explicit categorization task called the odd-man-out test, where participants were given three images among which two represented nodes next to each other on the network, and the remaining one was drawn from nodes at least three degrees away from the other two images. After completing the rotation task, participants were told that the stream of images they saw previously followed a pattern, and their goal in the odd-man-out test was to correctly select the image that did not fit with the other two images. If participants successfully learned the network structure, they would select the image that was far from the other two images on the network. Thus, an individual's network representation capacity could be quantified as her accuracy in the odd-man-out test.

The other behavioral indicator is participants' reaction time during the rotation task. If participants learned the underlying network structure, they were expected to spend longer on the rotation judgment task when the image to be judged came from a different cluster than the one they had previously seen. The rationale is that after learning the structure, they would anticipate the current image to come from the same cluster as the previous one. When the current image went against their expectation, it took longer to respond as they had to overcome the anticipation. This phenomenon is called surprisal effect (Schapiro et al. 2013). A larger surprisal effect (i.e., slower reaction time after cluster transitions relative to reaction time prior to cluster transitions) indicates better capacity to learn network representation.

Compared to structural learning, the statistical learning paradigm requires more trials for participants to acquire the network structure and more sophisticated design of outcome measures. However, it enables inferences of individual's global network representation capacity in a controlled experimental setting and is especially powerful for investigating how and under what circumstances, humans may switch between different global relational schemata to offload cognitive loads (Gebhart et al. 2009).

CONCLUSIONS

In this chapter, we reviewed the puzzle of people's low accuracy in perceptions of local ties in egocentric networks versus their much higher accuracy in perceptions of global network structures such as eigenvector centrality. We posit that such a puzzle is resolved by recognizing that humans rely on layers of relational schemata to mentally organize their social contacts. In other words, within-subject differences in local and global accuracy essentially reflect differences in the schemata used by the individual to mentally represent her social network information on the two levels. Across individuals, variations in the use of schemata can be attributed to individual differences in each person's schemata repertoire, as well as her tendencies to adopt certain schemata in a particular situation or context given prior experiences. Consequently, the specific set of schemata one activates could vary in its sufficiency and appropriateness in fully representing the network structure. This eventually introduces variations in individual capacity to mentally envision network characteristics from their mental representations.

Building upon this insight, we reviewed and compared four prominent survey or experimental paradigms that enable quantifications of individual differences in network representation capacity: the error paradigm, the free-recall paradigm, the structural learning paradigm and the statistical learning paradigm. While these four paradigms have different strengths and weaknesses (e.g., difficulty in administration) and are appropriate for examining network representation capacity on different levels (i.e., local network features or global network structures), they nonetheless invite a new research agenda that seeks to understand how individuals' mental organization of their social network information correspond with subjectively reported or objectively observed social connections.

Theoretically, our approach is consistent with the conceptual plurality of social networks as "methodological tools," as "metaphors for understanding forms of relations," and as "descriptors of social forms" (Knox, Savage, and Harvey 2006, p. 114). As methodological tools, social networks are logical models constructed by analysts to account for observed social structure in a coherent and economical way. Formulas that describe these models do not describe or prescribe any social practice and/or social process that generate the observed social structure. As metaphors, social networks symbolize social relatedness and interconnections. And as descriptors, social networks become referents to the social structure that they are designed to describe. In some sense, our question about individual mental representations of social networks connects the instrumentalists' toolkit view with the realists' descriptor view of social networks by investigating the micro-cognitive representation of social network structures in the human mind. When the "analysts" who develop models of social structure, including both researchers and those whom we study, also have the capacity to act and shape the structure they observe,

interaction between the two worlds – mental and social – becomes critical to understanding dynamics of social networks.

Although it remains an empirical question whether a perfect network representation is the best mental model for utilizing and mobilizing social network resources, there exists the possibility that envisaging and mobilizing social contacts in such a manner may eventually recreate social structure into social networks.

In the end, future research should explore the full cycle of egocentric networks from their source in social interactions, to their representation in each member's mind, and their eventual behavioral realization in practices of mobilization. Such efforts will enable researchers to better understand how network information is encoded into the human mind, and what relational schemata individuals adopt given their social connections. Additionally, we also encourage future researchers to investigate how and to what extent individuals incorporate the network diagram into their perceptions of social relationships. Realizing that network visualization itself can become a basis for social action, researchers may want to investigate the changing meanings of networks in the eyes of participants, especially as the rise of social media makes the concept of social networks more prevalent and increasingly subject to popular mythologizing.

REFERENCES

Alba, Joseph W., and Lynn Hasher. 1983. *Is Memory Schematic? Psychological Bulletin* 93(2): 203.

Albert, Réka, and Albert László Barabási. 2002. "Statistical Mechanics of Complex Networks." *Reviews of Modern Physics* 74(1): 47–97.

Barabási, Albert László, and Reka Albert. 1999. "Emergence of Scaling in Random Networks." *Science* 286(5439): 509–12.

Bartlett, Frederic Charles. 1932. *Remembering: A Study in Experimental and Social Psychology*. Cambridge: Cambridge University Press.

Berger, Jonah, and Chip Heath. 2008. "Who Drives Divergence? Identity Signaling, Outgroup Dissimilarity, and the Abandonment of Cultural Tastes." *Journal of Personality and Social Psychology* 95(3): 593–607.

Bernard, H. Russel., Peter D. Killworth, and Lee Sailer. 1982. "Informant Accuracy in Social-Network Data V. An Experimental Attempt to Predict Actual Communication from Recall Data." *Social Science Research* 11(1): 30–66.

Bond, Charles F., and Daniel R. Brockett. 1987. "A Social Context-Personality Index Theory of Memory for Acquaintances." *Journal of Personality and Social Psychology* 52(6): 1110–21.

Bond, Charles F., Rosalind L. Jones, and Daniel L. Weintraub. 1985. "On the Unconstrained Recall of Acquaintances: A Sampling-Traversal Model." *Journal of Personality and Social Psychology* 49(2): 327–37.

Borgatti, Stephen P., Kathleen M. Carley, and David Krackhardt. 2006. "On the Robustness of Centrality Measures under Conditions of Imperfect Data." *Social Networks* 28(2): 124–36.

Bousfield, W. A., and C. H. W. Sedgewick. 1944. "An Analysis of Sequences of Restricted Associative Responses." *Journal of General Psychology* 30(2): 149–65.

Brady, Timothy F., and Aude Oliva. 2008. "Statistical Learning Using Real-World Scenes." *Psychological Science* 19(7): 678–85.

Brashears, Matthew E. 2013. "Humans Use Compression Heuristics to Improve the Recall of Social Networks." *Scientific Reports* 3(1513): 1–7.

Brashears, Matthew E., and Eric Quintane. 2015. "The Microstructures of Network Recall: How Social Networks Are Encoded and Represented in Human Memory." *Social Networks* 41: 113–26.

Brashears, Matthew E., and Laura Aufderheide Brashears. 2016. "The Enemy of My Friend Is Easy to Remember: Balance as a Compression Heuristic," pp. 1–31 in *Advances in Group Processes*, edited by S. Thye and E. Lawler. Bingley: Emerald Insight.

Brashears, Matthew E., and Laura Aufderheide Brashears. 2020. "Compression Heuristics, Social Networks, and the Evolution of Human Intelligence." In *Network Science in Cognitive Psychology (Frontiers of Cognitive Psychology)*, edited by Michael S. Vitevitch. New York: Routledge.

Brewer, Devon D. 1995a. "Cognitive Indicators of Knowledge in Semantic Domains." *Journal of Quantitative Anthropology* 5(2): 107–28.

Brewer, Devon D. 1995b. "The Social Structural Basis of the Organization of Persons in Memory." *Human Nature* 6(4): 379–403.

Brewer, Devon D. 1997. "No Associative Biases in the First Name Cued Recall Procedure for Eliciting Personal Networks." *Social Networks* 19(4): 345–53.

Brewer, Devon D., Giovanni Rinaldi, Andrei Mogoutov, and Thomas W. Valente. 2005. "A Quantitative Review of Associative Patterns in the Recall of Persons." *Journal of Social Structure* 6(1).

Brewer, William F., and Glenn V. Nakamura. 1984. "The Nature and Functions of Schemas" in *Handbook of Social Cognition*, edited by R. S. Wyer and T. K. Srull. Hillsdale, NJ: Erlbaum.

Broido, Anna D., and Aaron Clauset. 2019. "Scale-Free Networks Are Rare." *Nature Communications* 10(1): 1017.

Burt, Ronald S. 2009. *Structural Holes: The Social Structure of Competition*. Cambridge, MA: Harvard University Press.

Burt, Ronald S., and Wm. M. Bittner. 1981. "A Note on Inferences Regarding Network Subgroups." *Social Networks* 3(1): 71–88.

Carlston, Donal E., and John J. Skowronski. 1994. "Savings in the Relearning of Trait Information as Evidence for Spontaneous Inference Generation." *Journal of Personality and Social Psychology* 66(5): 840–56.

Casciaro, Tiziana, Kathleen M. Carley, and David Krackhardt. 1999. "Positive Affectivity and Accuracy in Social Network Perception." *Motivation and Emotion* 23(4): 285–305.

Chambers, Kyle E., Kristine H. Onishi, and Cynthia Fisher. 2003. "Infants Learn Phonotactic Regularities from Brief Auditory Experience." *Cognition* 87(2): B69–B77.

Chase, William G., and Herbert A. Simon. 1973. "Perception in Chess." *Cognitive Psychology* 4(1): 55–81.

Chun, Marvin M., and Yuhong Jiang. 1999. "Top-Down Attentional Guidance Based on Implicit Learning of Visual Covariation." *Psychological Science* 10(4): 360–5.

Conway, Christopher M., and Morten H. Christiansen. 2005. "Modality-Constrained Statistical Learning of Tactile, Visual, and Auditory Sequences." *Journal of Experimental Psychology: Learning, Memory, and Cognition* 31(1): 24–39.

Crockett, Walter H. 1982. "Balance, Agreement, and Positivity in the Cognition of Small Social Structures." *Advances in Experimental Social Psychology* 15: 1–57.

Ebel, Holger, Lutz-Ingo Mielsch, and Stefan Bornholdt. 2002. "Scale-Free Topology of e-Mail Networks." *Physical Review E* 66(3): 035103.

Fiedler, Klaus. 1982. "Causal Schemata: Review and Criticism of Research on a Popular Construct." *Journal of Personality and Social Psychology* 42(6): 1001–13.

Fiori, Katherine. L., Jacqui Smith, and Toni C. Antonucci. 2007. "Social Network Types Among Older Adults: A Multidimensional Approach." *The Journals of Gerontology Series B: Psychological Sciences and Social Sciences* 62(6): P322–30.

Fischer, Claude S., Robert Max Jackson, C. Ann Stueve, Kathleen Gerson, Lynne McCallister Jones, and Mark Baldassare. 1977. *Networks and Places: Social Relations in the Urban Setting*. New York: Free Press.

Fiser, József, and Richard N. Aslin. 2001. "Unsupervised Statistical Learning of Higher-Order Spatial Structures from Visual Scenes." *Psychological Science* 12(6): 499–504.

Fiske, Alan Page. 1995. "Social Schemata for Remembering People: Relationships and Person Attributes in Free Recall of Acquaintances." *Journal of Quantitative Anthropology* 5: 305–24.

Fodor, Jerry A., and Zenon W. Pylyshyn. 1988. "Connectionism and Cognitive Architecture: A Critical Analysis." *Cognition* 28(1–2): 3–71.

Freeman, Linton C. 1992. "Filling in the Blanks: A Theory of Cognitive Categories and the Structure of Social Affiliation." *Social Psychology Quarterly* 55(2): 118–27.

Freeman, Linton C., A. Kimball Romney, and Sue C. Freeman. 1987. "Cognitive Structure and Informant Accuracy." *American Anthropologist* 89(2): 310–25.

Gebhart, Andrea L., Richard N. Aslin, and Elissa L. Newport. 2009. "Changing Structures in Midstream: Learning Along the Statistical Garden Path." *Cognitive Science* 33(6): 1087–116.

Gershman, Samuel J., Hillard Thomas Pouncy, and Hyowon Gweon. 2017. "Learning the Structure of Social Influence." *Cognitive Science* 41: 545–75.

Gobet, Fernand, and Herbert A. Simon. 1996. "Recall of Rapidly Presented Random Chess Positions Is a Function of Skill." *Psychonomic Bulletin & Review* 3(2):159–63.

Graesser, Arthur, and George Mandler. 1978. "Limited Processing Capacity Constrains the Storage of Unrelated Sets of Words and Retrieval from Natural Categories." *Journal of Experimental Psychology* 4(1): 86–100.

Granovetter, Mark S. 1973. "The Strength of Weak Ties." *American Journal of Sociology* 78(6): 1360–80.

Griffiths, Thomas L., Mark Steyvers, and Joshua B. Tenenbaum. 2007. "Topics in Semantic Representation." *Psychological Review* 114(2): 211–44.

Gruenewald, Paul J., and Gregory R. Lockhead. 1980. "The Free Recall of Category Examples." *Journal of Experimental Psychology* 6(3): 225–40.

Hamilton, David L., Von O. Leirer, and Lawrence B. Katz. 1979. "A Clustering Analysis of Organizational Processes in Impression Formation." Unpublished manuscript, University of California, Santa Barbara.

Harris, Richard J., and Gregory E. Monaco. 1978. "Psychology of Pragmatic Implication: Information Processing between the Lines." *Journal of Experimental Psychology* 107(1): 1–22.

Hill, R. A., and R. I. M. Dunbar. 2003. "Social Network Size in Humans." *Human Nature* 14(1): 53–72.

Hintzman, Douglas L. 1986. "'Schema Abstraction' in a Multiple-Trace Memory Model." *Psychological Review* 93(4): 411–28.

Holland, John H., Keith J. Holyoak, Richard E. Nisbett, and Paul R. Thagard. 1989. *Induction: Processes of Inference, Learning, and Discovery*. Cambridge, MA: MIT Press.

Holyoak, Keith J. 2008. "Induction as Model Selection." *Proceedings of the National Academy of Sciences of the United States of America* 105(31): 10637–8.

Janicik, Gregory A., and Richard P. Larrick. 2005. "Social Network Schemas and the Learning of Incomplete Networks." *Journal of Personality and Social Psychology* 88 (2): 348–64.

John Scott. 1988. "Social Network Analysis." *Sociology* 22(1): 109–27.

Karuza, Elisabeth A., Sharon L. Thompson-Schill, and Danielle S. Bassett. 2016. "Local Patterns to Global Architectures: Influences of Network Topology on Human Learning." *Trends in Cognitive Sciences* 20(8): 629–40.

Kelley, Harold H. 1967. "Attribution Theory in Social Psychology." In *Nebraska Symposium on Motivation*, vol. 15, edited by D. Levine. Lincoln: University of Nebraska Press.

Kemp, Charles, and Joshua B. Tenenbaum. 2008. "The Discovery of Structural Form." *Proceedings of the National Academy of Sciences of the United States of America* 105 (31): 10687–92.

Kidd, Celeste, Steven T. Piantadosi, and Richard N. Aslin. 2012. "The Goldilocks Effect: Human Infants Allocate Attention to Visual Sequences That Are Neither Too Simple Nor Too Complex" edited by A. Rodriguez-Fornells. *PLoS ONE* 7(5): e36399.

Kim, Robyn, Aaron Seitz, Heather Feenstra, and Ladan Shams. 2009. "Testing Assumptions of Statistical Learning: Is It Long-Term and Implicit?" *Neuroscience Letters* 461(2): 145–9.

Kimchi, Ruth. 2003. "Relative Dominance of Holistic and Component Properties in the Perceptual Organization of Visual Objects," pp. 235–68 in *Perception of Faces, Objects, and Scenes: Analytic and Holistic Processes*, edited by M. A. Peterson and G. Rhodes. New York: Oxford University Press.

Kirkham, Natasha Z., Jonathan A. Slemmer, and Scott P. Johnson. 2002. "Visual Statistical Learning in Infancy: Evidence for a Domain General Learning Mechanism." *Cognition* 83(2): B35–42.

Knox, Hannah, Mike Savage, and Penny Harvey. 2006. "Social Networks and the Study of Relations: Networks as Method, Metaphor and Form." *Economy and Society* 35 (1): 113–40.

Krackhardt, David. 1987. "Cognitive Social Structures." *Social Networks* 9: 109–34.

Krackhardt, David. 1996. "Comment on Burt and Knez's Third-Party Effects on Trust." *Rationality and Society* 8(1): 111–20.

Krackhardt, David. 2014. "A Preliminary Look at Accuracy in Egonets." *Research in the Sociology of Organizations* 40(2014): 277–93.

Krackhardt, David, and Martin Kilduff. 1999. "Whether Close or Far: Social Distance Effects on Perceived Balance in Friendship Networks." *Journal of Personality and Social Psychology* 76(5): 770–82.

Kumbasar, Ece, A. Kimball Rommey, and William H. Batchelder. 1994. "Systematic Biases in Social Perception." *American Journal of Sociology* 100(2): 477–505.

Lansu, Tessa A. M., Antonius H. N. Cillessen, and Johan C. Karremans. 2014. "Adolescents' Selective Visual Attention for High-Status Peers: The Role of Perceiver Status and Gender." *Child Development* 85(2): 421–8.

Lin, Nan. 1999. "Social Networks and Status Attainment." *Annual Review of Sociology* 25(1): 467–87.

Markman, Arthur B. 1999. *Knowledge Representation*. Mahwah, NJ: Lawrence Erlbaum Associates Publishers.

Medin, Douglas L., Elizabeth B. Lynch, John D. Coley, and Scott Atran. 1997. "Categorization and Reasoning among Tree Experts: Do All Roads Lead to Rome?" *Cognitive Psychology* 32(1): 49–96.

Messinger, A., L. R. Squire, S. M. Zola, and T. D. Albright. 2001. "Neuronal Representations of Stimulus Associations Develop in the Temporal Lobe during Learning." *Proceedings of the National Academy of Sciences of the United States of America* 98(21): 12239–44.

Minsky, Marvin. 1975. "A Framework for Representing Knowledge." in *The Psychology of Computer Vision*, edited by P. H. Winston. New York: McGraw-Hill.

Newell, Allen, and Herbert Alexander Simon. 1972. *Human Problem Solving*. Englewood Cliffs, NJ: Prentice-Hall.

Newman, M. E. J., and Juyong Park. 2003. "Why Social Networks Are Different from Other Types of Networks." *Physical Review E* 68(3): 036122.

Onnis, Luca, Heidi R. Waterfall, and Shimon Edelman. 2008. "Learn Locally, Act Globally: Learning Language from Variation Set Cues." *Cognition* 109(3): 423–30.

Osada, Takahiro, Yusuke Adachi, Hiroko M. Kimura, and Yasushi Miyashita. 2008. "Towards Understanding of the Cortical Network Underlying Associative Memory." *Philosophical Transactions of the Royal Society B: Biological Sciences* 363(1500): 2187–99.

Paivio, Allan. 1969. "Mental Imagery in Associative Learning and Memory." *Psychological Review* 76(3): 241–63.

Parkinson, Carolyn, Adam M. Kleinbaum, and Thalia Wheatley. 2017. "Spontaneous Neural Encoding of Social Network Position." *Nature Human Behaviour* 1(5): 72.

Pescosolido, Bernice A., and Eric R. Wright. 2004. "The View from Two Worlds: The Convergence of Social Network Reports between Mental Health Clients and Their Ties." *Social Science and Medicine* 58(9): 1795–806.

Prentice, Deborah A. 1990. "Familiarity and Differences in Self- and Other-Representations." *Journal of Personality and Social Psychology* 59(3): 369–83.

Roenker, Daniel L., Charles P. Thompson, and Sam C. Brown. 1971. "Comparison of Measures for the Estimation of Clustering in Free Recall." *Psychological Bulletin* 76 (1): 45–8.

Rumelhart, David E., and James L. Mcclelland. 1986. *Parallel Distributed Processing*. Cambridge, MA: MIT Press.

Rumelhart, David E., and Andrew Ortony. 1978. "The Representation of Knowledge in Memory," pp. 99–135 in *Schooling and the Acquisition of Knowledge*, edited by R. C. Anderson, R. J. Spiro, and W. E. Montague. Oxford: Lawrence Erlbaum.

Saariluoma, Pertti. 1989. "Chess Players' Recall of Auditorily Presented Chess Positions." *European Journal of Cognitive Psychology* 1(4): 309–20.

Saffran, Jenny R., Richard N. Aslin, and Elissa L. Newport. 1996. "Statistical Learning by 8-Month-Old Infants." *Science* 274(5294): 1926–8.

Saffran, Jenny R., and Diana P. Wilson. 2003. "From Syllables to Syntax: Multilevel Statistical Learning by 12-Month-Old Infants." *Infancy* 4(2): 273–84.

Schank, Roger C., and Robert P. Abelson. 1977. *Scripts, Plans, Goals and Understanding: An Inquiry into Human Knowledge Structures*. Oxford: Lawrence Erlbaum.

Schapiro, Anna C., and Nicholas B. Turk-Browne. 2015. "Statistical Learning," pp. 501–6 in *Brain Mapping: An Encyclopedic Reference*, edited by A. W. Toga. Cambridge, MA: Elsevier.

Schapiro, Anna C., Emma Gregory, Barbara Landau, Michael McCloskey, and Nicholas B. Turk-Browne. 2014. "The Necessity of the Medial Temporal Lobe for Statistical Learning." *Journal of Cognitive Neuroscience* 26(8): 1736–47.

Schapiro, Anna C., Timothy T. Rogers, Natalia I. Cordova, Nicholas B. Turk-Browne, and Matthew M. Botvinick. 2013. "Neural Representations of Events Arise from Temporal Community Structure." *Nature Neuroscience* 16(4): 486–92.

Smolensky, Paul. 1988. "On the Proper Treatment of Connectionism." *Behavioral and Brain Sciences* 11(1): 1–23.

De Soto, Clinton B. 1960. "Learning a Social Structure." *Journal of Abnormal and Social Psychology* 60(3): 417–21.

De Soto, Clinton B., Nancy M. Henley, and Marvin London. 1968. "Balance and the Grouping Schema." *Journal of Personality and Social Psychology* 8(1, Pt.1): 1–7.

Sulin, Rebecca A., and D. James Dooling. 1974. "Intrusion of a Thematic Idea in Retention of Prose." *Journal of Experimental Psychology* 103(2): 255–62.

Sun, Hui, and Edward B. Smith. 2018. "Cognitive Search and Social Networks." Academy of Management Conference, Chicago, United States.

Tang, Jian, Meng Qu, Mingzhe Wang, Ming Zhang, Jun Yan, and Qiaozhu Mei. 2015. "LINE: Large-Scale Information Network Embedding," pp. 1067–77 in the *Proceedings of the 24th International Conference on the World Wide Web*.

Tavares, Rita Morais, Avi Mendelsohn, Yael Grossman, Christian Hamilton Williams, Matthew Shapiro, Yaacov Trope, and Daniela Schiller. 2015. "A Map for Social Navigation in the Human Brain." *Neuron* 87(1): 231–43.

Todorov, Alexander, and James S. Uleman. 2003. "The Efficiency of Binding Spontaneous Trait Inferences to Actors' Faces." *Journal of Experimental Social Psychology* 39(6): 549–62.

Tompson, Steven, Ari Kahn, Emily Falk, Jean Vettel, and Danielle Bassett. 2019. "Individual Differences in Learning Social and Non-Social Network Structures." *Journal of Experimental Psychology* 45(2): 253–71.

Treisman, Anne. 1986. "Features and Objects in Visual Processing." *Scientific American* 255(5): 114–25.

Uleman, James S., S. Adil Saribay, and Celia M. Gonzalez. 2007. "Spontaneous Inferences, Implicit Impressions, and Implicit Theories." *Annual Review of Psychology* 59: 329–60.

Vapnik, Vladimir N. 1999. *An Overview of Statistical Learning Theory. IEEE Transactions on Neural Networks* 10(5): 988–99.

Vaughn, Brian E., and Everett Waters. 1981. "Attention Structure, Sociometric Status, and Dominance: Interrelations, Behavioral Correlates, and Relationships to Social Competence." *Developmental Psychology* 17(3): 275–88.

Wagemans, Johan, Jacob Feldman, Sergei Gepshtein, Ruth Kimchi, James R. Pomerantz, Peter A. van der Helm, and Cees van Leeuwen. 2012. "A Century of Gestalt Psychology in Visual Perception: II. Conceptual and Theoretical Foundations." *Psychological Bulletin* 138(6): 1218–52.

Watts, Duncan J., and Steven H. Strogatz. 1998. "Collective Dynamics of 'Small-World' Networks." *Nature* 393(6684): 440–2.

Winkler-Rhoades, Nathan, Douglas Medin, Sandra R. Waxman, Jennie Woodring, and Norbert O. Ross. 2010. "Naming the Animals That Come to Mind: Effects of Culture and Experience on Category Fluency." *Journal of Cognition and Culture* 10(1–2): 205–20.

Wood, Gordon. 1978. "The Knew-It-All-along Effect." *Journal of Experimental Psychology* 4(2): 345–53.

Yonelinas, Andrew P. 1994. "Receiver-Operating Characteristics in Recognition Memory: Evidence for a Dual-Process Model." *Journal of Experimental Psychology* 20(6): 1341–54.

Yonelinas, Andrew P. 2002. "The Nature of Recollection and Familiarity: A Review of 30 Years of Research." *Journal of Memory and Language* 46: 441–517.

Zerubavel, Noam, Peter S. Bearman, Jochen Weber, and Kevin N. Ochsner. 2015. "Neural Mechanisms Tracking Popularity in Real-World Social Networks." *Proceedings of the National Academy of Sciences of the United States of America* 112(49): 15072–7.

Zhou, W. X., D. Sornette, R. A. Hill, and R. I. M. Dunbar. 2005. "Discrete Hierarchical Organization of Social Group Sizes." *Proceedings. Biological Sciences* 272(1561): 439–44.