



How humans learn and represent networks

Christopher W. Lynn^a and Danielle S. Bassett^{a,b,c,d,e,f,1} 

^aDepartment of Physics & Astronomy, College of Arts & Sciences, University of Pennsylvania, Philadelphia, PA 19104; ^bDepartment of Bioengineering, School of Engineering & Applied Science, University of Pennsylvania, Philadelphia, PA 19104; ^cDepartment of Electrical & Systems Engineering, School of Engineering & Applied Science, University of Pennsylvania, Philadelphia, PA 19104; ^dDepartment of Neurology, Perelman School of Medicine, University of Pennsylvania, Philadelphia, PA 19104; ^eDepartment of Psychiatry, Perelman School of Medicine, University of Pennsylvania, Philadelphia, PA 19104; and ^fSanta Fe Institute, Santa Fe, NM 87501

Edited by Richard M. Shiffrin, Indiana University, Bloomington, IN, and approved February 13, 2020 (received for review September 15, 2019)

Humans receive information from the world around them in sequences of discrete items—from words in language or notes in music to abstract concepts in books and websites on the Internet. To model their environment, from a young age people are tasked with learning the network structures formed by these items (nodes) and the connections between them (edges). But how do humans uncover the large-scale structures of networks when they experience only sequences of individual items? Moreover, what do people’s internal maps and models of these networks look like? Here, we introduce graph learning, a growing and interdisciplinary field studying how humans learn and represent networks in the world around them. Specifically, we review progress toward understanding how people uncover the complex webs of relationships underlying sequences of items. We begin by describing established results showing that humans can detect fine-scale network structure, such as variations in the probabilities of transitions between items. We next present recent experiments that directly control for differences in transition probabilities, demonstrating that human behavior depends critically on the mesoscale and macroscale properties of networks. Finally, we introduce computational models of human graph learning that make testable predictions about the impact of network structure on people’s behavior and cognition. Throughout, we highlight open questions in the study of graph learning that will require creative insights from cognitive scientists and network scientists alike.

graph learning | cognitive science | network science | statistical learning | knowledge networks

Our experience of the world is punctuated by discrete items and events, all connected by a hidden network of forces, causes, and associations. Just as navigation requires a mental map of one’s physical surroundings (1, 2), anticipation, planning, perception, and communication all depend on people’s ability to learn the network structure connecting items and events in their environment (3–5). For example, to identify the boundaries between words, children as young as 8 mo old identify subtle variations in the network of transitions between syllables in spoken language (6). Within their first 30 mo, toddlers already learn enough words to form complex language networks that exhibit robust structural features (7–9). By the time we reach adulthood, graph learning enables us to understand and produce language (6, 10), flexibly and adaptively learn words (11, 12), parse continuous streams of stimuli (6), build social intuitions (13), perform abstract reasoning (14), and categorize visual patterns (15). In this way, our ability to learn the structures of networks supports a wide range of cognitive functions.

Our capacity to infer and represent complex relationships has also enabled humans to construct an impressive array of networked systems, from language (16–18) and music (19) to social networks (20, 21), the Internet (22, 23), and the web of concepts that constitute the arts and sciences (24, 25). Moreover, individual differences in cognition, such as those driven by learning disabilities and age, give rise to variations in the types of network structures that people are able to construct (26, 27). Therefore, studying how humans learn and represent

networks will not only inform our understanding of how we perform many of our basic cognitive functions, but also shed light on the structure and function of networks in the world around us.

Here, we provide a brief introduction to the field of graph learning, spanning the experimental techniques and network-based models, theories, and intuitions recently developed to study the effects of network structure on human cognition and behavior. Given the highly interdisciplinary nature of the field—which draws upon experimental methods from cognitive science and linguistics and builds upon computational techniques from network science, information theory, and statistical learning—we aim to present an accessible overview with simple motivating examples.

We focus particular attention on understanding how people uncover the structure of connections between items in a sequence, such as syllables and words in spoken and written language, concepts in books and classroom lectures, or notes in musical progressions. We begin by discussing experimental results demonstrating that humans are adept at detecting differences in the probabilities of transitions between items and how such transitions connect and combine to form networks that encode the large-scale structure of entire sequences. We then present recent experiments that measure the effects of network structure on human behavior by directly controlling for differences in transition probabilities, followed by a description of the computational models that have been proposed to account for these network effects. We conclude by highlighting some of the open research directions stemming from recent advances in graph learning, including important generalizations of existing graph learning paradigms and direct implications for understanding the structure and function of real-world networks.

Learning Transition Probabilities

As humans navigate their environment and accumulate experience, one of the brain’s primary functions is to infer the statistical relationships governing causes and effects (28, 29). Given a sequence of items, perhaps the simplest statistics available

This paper results from the Arthur M. Sackler Colloquium of the National Academy of Sciences, “Brain Produces Mind by Modeling,” held May 1–3, 2019, at the Arnold and Mabel Beckman Center of the National Academies of Sciences and Engineering in Irvine, CA. NAS colloquia began in 1991 and have been published in PNAS since 1995. From February 2001 through May 2019, colloquia were supported by a generous gift from The Dame Jillian and Dr. Arthur M. Sackler Foundation for the Arts, Sciences, & Humanities, in memory of Dame Sackler’s husband, Arthur M. Sackler. The complete program and video recordings of most presentations are available on the NAS website at <http://www.nasonline.org/brain-produces-mind-by>.

Author contributions: C.W.L. and D.S.B. designed research; and C.W.L. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

Published under the [PNAS license](#).

¹To whom correspondence may be addressed. Email: dsb@seas.upenn.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.1912328117/-DCSupplemental>.

First published November 23, 2020.

to a learner are the frequencies of transitions from one item to another. Naturally, the field of statistical learning, which is devoted to understanding how humans extract statistical regularities from their environment, has predominantly focused on these simple statistics. For example, consider spoken language, wherein distinct syllables transition from one to another in a continuous stream without pauses or demarcations between words (30). How do people segment such continuous streams of data, identifying where one word starts and another begins? The answer, as research has robustly established (31–34), lies in the statistical properties of the transitions between syllables.

The ability to detect words within continuous speech was initially demonstrated by Saffran et al. (6), who exposed infants to sequences of four pseudowords, each consisting of three syllables (Fig. 1A). The order of syllables within each word remained consistent, yielding a within-word transition probability of 1. However, the order of the words was random, yielding a between-word transition probability of 1/3. Infants were able to reliably detect this difference in syllable transition probabilities, thereby providing a compelling mechanism for word identification during language acquisition. This experimental paradigm has since been generalized to study statistical learning in other domains, with stimuli ranging from colors (35) and shapes (15) to visual scenes (36) and physical actions (37). Indeed, the capacity to uncover variations in transition probabilities is now recognized as a central and general feature of human learning (31–34).

Learning Network Structure

Although individual connections between items provide important information about the structure of a system, they do not

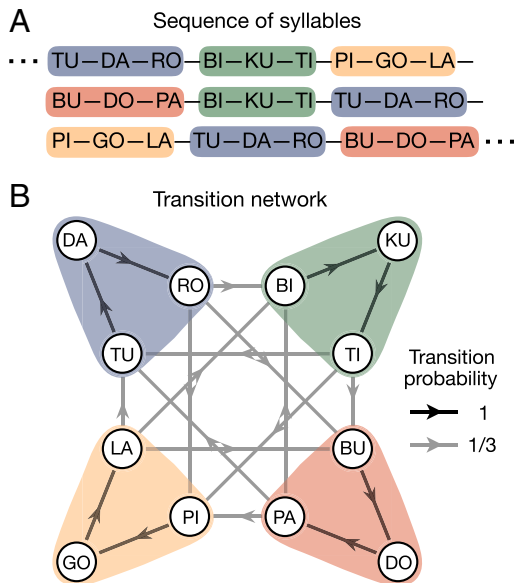


Fig. 1. Transitions between syllables in the fabricated language of Saffran et al. (6). (A) A sequence containing four different pseudowords: *tudaro* (blue), *bikuti* (green), *budopa* (red), and *pigola* (yellow). When spoken, the sequence forms a continuous stream of syllables, without clear boundaries between words. The transition probability from one syllable to another is 1 if the transition occurs within a word and 1/3 if the transition occurs between words. This difference in transition probabilities allows infants to segment spoken language into distinct words (6, 31, 38). (B) The transitions between syllables form a network, with edge weights representing the syllable transition probabilities. A random walk in the transition network defines a sequence of syllables in the pseudolanguage. The four pseudowords form distinct communities (highlighted regions) that are easily identifiable by eye. Adapted from ref. 38, with permission from Elsevier.

tell the whole story. Connections also combine and overlap to form complex webs that characterize the higher-order structure of our environment. To study these structures, scientists have increasingly turned to the language of network science (40), conceptualizing items as nodes in a network with edges defining possible connections between them (see *SI Appendix, Fig. S1* for a primer on networks). One can then represent a sequence of items, such as the stream of syllables in spoken language, as a walk through this underlying network (19, 41–44). This perspective has been particularly useful in the study of artificial grammar learning (45–47), wherein human subjects are tasked with inferring the grammar rules (i.e., the network of transitions between letters and words) underlying a fabricated language.

By translating items and connections into the language of network science, one inherits a powerful set of descriptive tools and visualization techniques for characterizing different types of structures. For example, consider once again the statistical learning experiment of Saffran et al. (ref. 6 and Fig. 1A). Simply by visualizing the transition structure as a network (Fig. 1B), it becomes clear that the syllables split naturally into four distinct clusters corresponding to the four different words in the artificial language. This observation raises an important question: When parsing words (or performing any other learning task), are people sensitive only to differences in individual connections, or do they also uncover large-scale features of the underlying network? In what follows, we describe recent advances in graph learning that shed light on precisely this question.

Learning Local Structure. The simplest properties of a network are those corresponding to individual nodes and edges, such as the weight of an edge, which determines the strength of the connection between two nodes, and the degree of a node, or its number of connections. For example, edge weights can represent transition probabilities between syllables or words (31–34), similarities between different semantic concepts (5, 16), or strengths of social interactions (20, 21). Meanwhile, significant effort has focused on understanding how humans learn the network structure surrounding individual nodes (8, 48–53). For example, the degree defines the connectedness of a node, such as the number of links pointing to a website (22, 23, 54), the number of friends that a person has (20), or the number of citations accumulated by a scientific paper (25). Notably, many of the networks that people encounter on a daily basis—including language, social, and hyperlink networks—exhibit heavy-tailed degree distributions, with many nodes of low degree and a select number of high-degree hubs (5, 16–18, 22, 24, 54–56).

Significant research has now demonstrated that people are able to learn the local network properties of individual nodes and edges, such as the transition probabilities between syllables in the previous section (31–34). To illustrate the impact of network structure on human behavior, we consider a recently developed experimental paradigm (42, 43), while noting that similar results have also been achieved using variations on this approach (13, 38, 39, 41, 44, 57). Specifically, each subject is shown a sequence of stimuli, with the order of stimuli defined by a random walk on an underlying transition network (Fig. 2A). Subjects are asked to respond to each stimulus by performing an action (and to avoid confounds the assignment of stimuli to nodes in the network is randomized across subjects). By measuring the speed with which subjects respond to stimuli, one can infer their expectations about the network structure: A fast reaction reflects a strongly anticipated transition, while a slow reaction reflects a weakly anticipated (or surprising) transition (42, 43, 58, 59).

Intuitively, one should expect a subject's anticipation to increase (and thus the reaction time to decrease) for edges representing more probable transitions. To test this prediction, we note that for a random walk in an unweighted and undirected

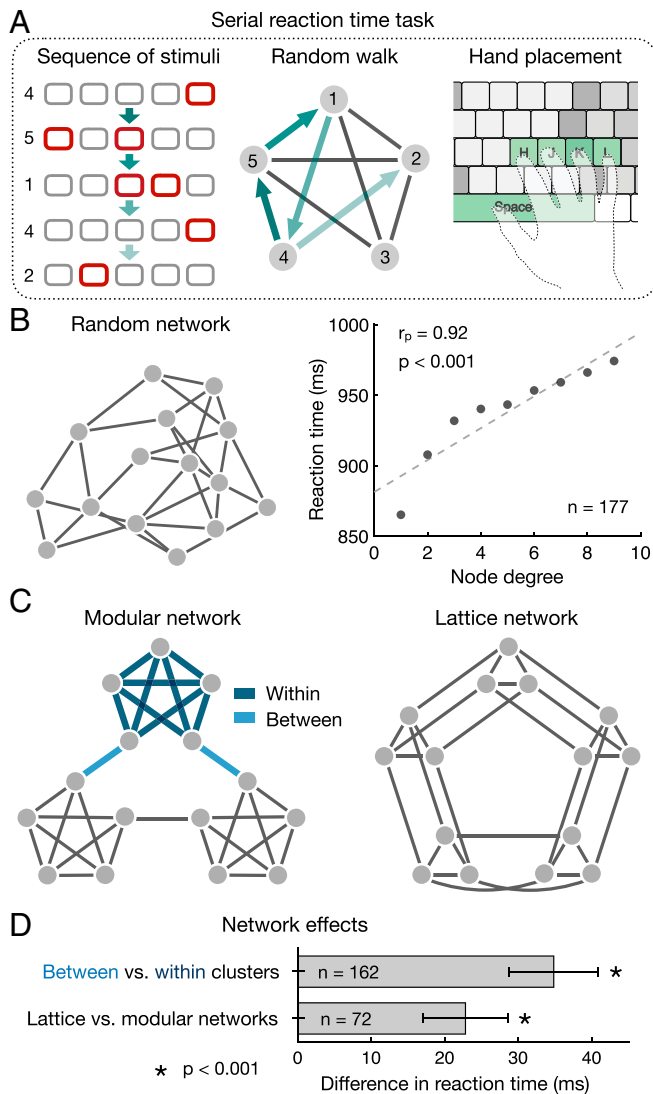


Fig. 2. Human behavior depends on network topology. (A) We consider a serial reaction time experiment in which subjects are shown sequences of stimuli and are asked to respond by performing an action. Here, each stimulus consists of five squares, one or two of which are highlighted in red (Left); the order of stimuli is determined by a random walk on an underlying network (Center); and for each stimulus, the subject presses the keys on the keyboard corresponding to the highlighted squares (Right). (B) Considering Erdős-Rényi random transition networks with 15 nodes and 30 edges (Left), subjects' average reaction times to a transition $i \rightarrow j$ increase as the degree k_i of the preceding node increases (Right). Equivalently, subjects' reaction times increase as the transition probability $P_{ij} = 1/k_i$ decreases (43). (C) To control for variations in transition probabilities, we consider two networks with constant degree $k = 4$: a modular network consisting of three communities of five nodes each (Left) and a lattice network representing a 3×5 grid with periodic boundary conditions (Right). (D) Experiments indicate two consistent effects of network structure. First, in the modular network, reaction times for between-cluster transitions are longer than for within-cluster transitions (39, 42, 43, 57). Second, reaction times are longer on average for the lattice network than for the modular network (42, 43).

network, the transition probability from one node i to a neighboring node j is given by $P_{ij} = 1/k_i$, where k_i is the degree of node i . Aligning with intuition, researchers have shown that people's reaction times are positively correlated with the degree of the previous stimulus (Fig. 2B), and therefore people are better able to anticipate more probable transitions (42, 43). Interestingly, significant research has also established similar results in

language networks, with people reading words more quickly if they occur more frequently or appear in more contexts (48, 49, 60). Conversely, humans tend to slow down and produce more errors when attempting to recall words with a large number of semantic associations, a phenomenon known as the fan effect (61, 62). Together, these results demonstrate that humans are sensitive to variations in the local properties of individual nodes and edges, but what about the mesoscale and macroscale properties of a network?

Learning Mesoscale Structure. The mesoscale structure of a network reflects the organizational properties of groups of nodes and edges. One such property is clustering or the tendency for a pair of nodes with a common neighbor to form a connection themselves. This tendency is clearly observed in social networks, where people with a common friend are themselves more likely to become friends. Similar principles govern the mesoscale structure of many other real-world networks, with items such as words, scientific papers, and webpages all exhibiting high clustering (25, 63–65). As nodes cluster together, they often give rise to a second mesoscale property—modular structure—which is characterized by tightly connected modules or communities of nodes. Such modular structure is now recognized as a ubiquitous feature of networks in our environment (66), with language splitting into groups of semantically or phonetically similar words (14, 18), people forming social cliques (20, 21, 67), and websites clustering into online communities (22).

Over the past 10 y, researchers have made significant strides toward understanding how the mesoscale properties of a network impact human learning and behavior. Words with higher clustering are more likely to be acquired during language learning (52), while words with lower clustering are easier to recognize in long-term memory (68) and convey processing (51, 53) and production (69) benefits. Additionally, in a series of cognitive and neuroimaging experiments, researchers have found that a network's modular structure has a significant impact on human behavior and neural activity. For example, people are able to detect the boundaries between communities in a network just by observing sequences of nodes (39, 41–43, 57). Moreover, strong modular structure helps people build more accurate mental representations of a network, thereby allowing humans to better anticipate future items and events (39, 41–43, 57).

Learning Global Structure. In addition to their local and mesoscale features, networks also have global properties that depend on the entire architecture of nodes and edges. Perhaps the most well-studied global property is small-world structure, wherein each node connects to every other node in only a small number of steps (65). Small-world topology has been observed in an array of networks that humans are tasked with learning, including social relationships (70), web hyperlinks (23), scientific citations (25), and semantic associations in language (18, 56). Moreover, in a particularly compelling example of the relationship between global network structure and human cognition, the small-world structure of people's learned language networks has been shown to vary from person to person, decreasing with age (27) and in people with learning disabilities (26).

While small-worldness describes the structure of an entire network, there are also measures that relate individual nodes to a network's global topology, including centrality (a measure of a node's role in mediating long-distance connections), communicability (a measure of the number of paths connecting a pair of nodes), and coreness (a measure of how deeply embedded a node is in a network). Global measures such as these have recently been shown to impact human learning and cognition, indicating that humans are sensitive to the global structure of networks in their environment. For example, in the reaction time experiments described above (Fig. 2A), people responded

more quickly and therefore were better able to anticipate nodes with low centrality (42). In a related experiment, neural activity was shown to reflect the communicability between pairs of stimuli in an underlying transition network (44). Finally, as children learn language, they more readily acquire and produce words with low coreness (50). Together, these results point to a robust and general relationship between large-scale network structure and human cognition. However, might these large-scale network effects simply be driven by confounding variations in the local network structure?

Controlling for Differences in Local Structure. To disentangle the effects of large-scale network structure from those of local structure, recent research has directly controlled for differences in transition probabilities by focusing on specific families of networks (39, 41–43). Recall that for random walks on unweighted, undirected networks, the transition probabilities are determined by node degrees. Therefore, to ensure that all transitions have equal probability, one can simply focus on graphs with constant degree but varying topology. For example, consider the modular and lattice graphs shown in Fig. 2C. Since both networks have constant degree 4 (and therefore constant transition probability 1/4 across all edges), any variation in behavior or cognition between different parts of a network, or between the two networks themselves, must stem from the networks' global topologies.

This approach was first developed by Schapiro et al. (41), who demonstrated that people are able to detect the transitions between clusters in the modular graph (Fig. 2C) and that these between-cluster transitions yield distinct patterns of neural activity relative to within-cluster transitions. Returning to the reaction time experiment (Fig. 2A), it was shown that subjects react more quickly to (and therefore are able to better anticipate) within-cluster transitions than between-cluster transitions (42, 43) (Fig. 2D). Moreover, people exhibit an overall decrease in reaction times for the modular graph relative to the lattice graph (42, 43) (Fig. 2D).

These results, combined with findings in similar experiments (39, 57), demonstrate that humans are sensitive to features of mesoscale and global network topology, even after controlling for differences in local structure. Thus, not only are humans able to learn individual transition probabilities, as originally demonstrated in seminal statistical learning experiments (Fig. 1), they are also capable of uncovering some of the complex structures found in our environment. But how do people learn the large-scale features of networks from past observations?

Modeling Human Graph Learning

Experiments spanning cognitive science, neuroscience, linguistics, and statistical learning have established that human behavior and cognition depend on the mesoscale and global topologies of networks in their environment. To understand how people detect these global features, and to make quantitative predictions about human behavior, one requires computational models of how humans construct internal representations of networks from past experiences. Here, we again focus on understanding how people learn the networks of transitions underlying observed sequences of items, such as words in a sentence, concepts in a book or classroom lecture, or notes in a musical progression. Interestingly, humans systematically deviate from the most accurate, and perhaps the simplest, learning rule.

To make these ideas concrete, consider a sequence of items described by the transition probability matrix P , where P_{ij} represents the conditional probability of one item i transitioning to another item j . Given an observed sequence of items, one can imagine estimating P_{ij} by simply dividing the number of times i has transitioned to j (denoted by n_{ij}) by the number of times i has appeared (which equals $\sum_k n_{ik}$):

$$\hat{P}_{ij} = \frac{n_{ij}}{\sum_k n_{ik}}. \quad [1]$$

In fact, not only is this perhaps the simplest estimate one could perform, it is also the most accurate (or maximum-likelihood) estimate of the transition probabilities from past observations (71). An important feature of maximum-likelihood estimation is that it gives an unbiased approximation of the true transition probabilities; that is, the estimated transition probabilities \hat{P}_{ij} are evenly distributed about their true values P_{ij} , independent of the large-scale structure of the network (71). However, we have seen that people's behavior and cognition depend systematically on mesoscale and global network properties, even when

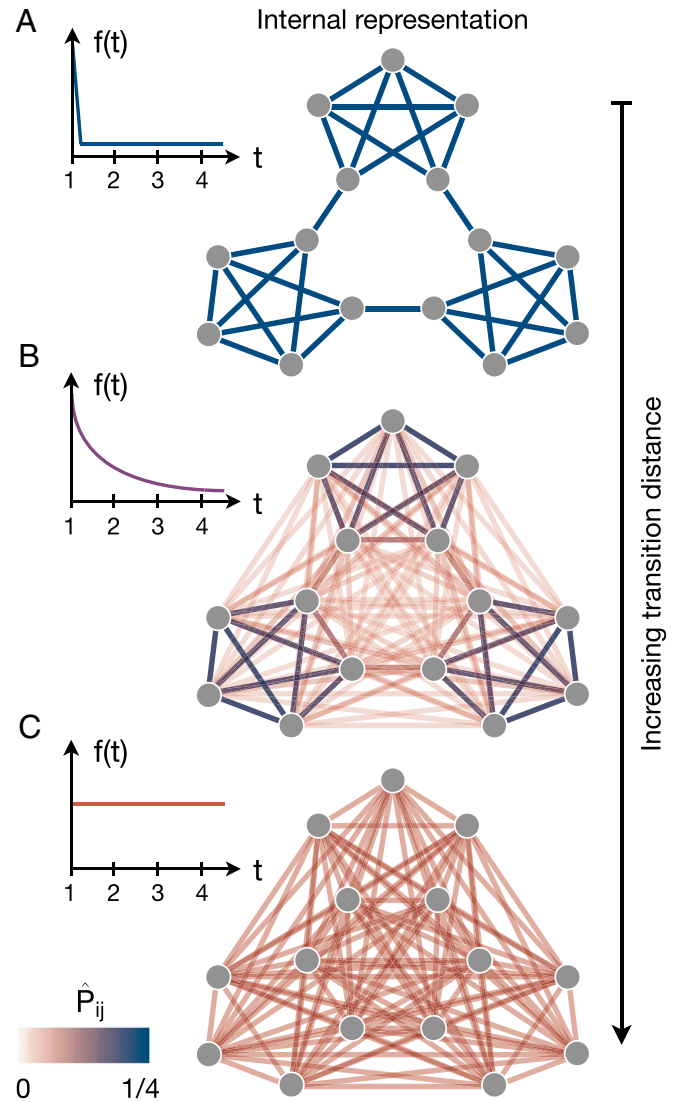


Fig. 3. Mesoscale and global network features emerge from long-distance associations. (A) Illustration of the weight function $f(t)$ (Left) and the learned network representation \hat{P} for learners that consider only transitions of length one. The estimated structure resembles the true modular network. (B) For learners that down-weight transitions of longer distances, higher-order features of the transition network, such as community structure, organically come into focus, yielding higher expected probabilities for within-cluster transitions than for between-cluster transitions. (C) For learners that equally weigh transitions of all distances, the internal representation becomes all to all, losing any resemblance to the true transition network. A–C correspond to learners that include progressively longer transitions in their network estimates. Adapted with permission from ref. 43.

transition probabilities are held constant (39, 41–44). Thus, when constructing internal representations, humans allow higher-order network structure to influence their estimates of individual transition probabilities, thereby deviating from maximum-likelihood estimation (43).

To understand the impact of network topology on human cognition, researchers have recently proposed a number of models describing how humans learn and represent transition networks (43, 44, 72–77). Notably, many of these models share a common underlying mechanism: that instead of just counting transitions of length one (as in maximum-likelihood estimation), humans also include transitions of lengths two, three, or more in their representations (43, 44, 75–78). Mathematically, by combining transitions of different distances, the estimated transition probabilities take the form

$$\hat{P}_{ij} = C \sum_{t \geq 1} f(t) n_{ij}^{(t)}, \quad [2]$$

where $n_{ij}^{(t)}$ represents the number of times that i has transitioned to j in t steps, $f(t)$ defines the weight placed on transitions of a

given distance, and C is a normalization constant. Interestingly, this simple prediction can be derived from a number of different cognitive theories—including the temporal context model of episodic memory (72), temporal difference learning and the successor representation in reinforcement learning (79–81), and the free energy principle from information theory (43). But how does combining transitions over different distances allow people to learn the structure of a network?

To answer this question, it helps to consider different choices for the function $f(t)$. Typically, $f(t)$ is assumed to be decreasing such that longer-distance associations contribute more weakly to a person's network representation (43, 79, 81). If $f(t)$ is a delta function centered at $t = 1$ (Fig. 3A), then the learner focuses on transitions of length one. In this case, people simply perform maximum-likelihood estimation, resulting in an unbiased estimate of the true transition structure P . Conversely, if $f(t)$ is uniform over all timescales $t \geq 1$, then the learner equally weighs transitions of all distances (Fig. 3C), and the estimate \hat{P} loses any resemblance to the true transition structure P . Importantly, however, for learners who combine transitions over intermediate distances (Fig. 3B), we find that large-scale features of




Fig. 4. Generalizations of the graph learning paradigm. (A) Transition networks often shift and change over time. Such nonstationary transition probabilities can be described using dynamical transition networks, which evolve from one network (for example, the modular network at *Left*) to another (for example, the ring network at *Right*) by iteratively rewiring edges. (B) Many real-world sequences have long-range dependencies, such that the next state depends not just on the current state, but also on a number of previous states (89, 90). For example, path 1 in the displayed network yields two possibilities for the next state (*Left*), while path 2 yields a different set of three possible states (*Right*). (C) Humans often actively seek out information by choosing their path through a transition network, rather than simply being presented with a prescribed sequence. Such information seeking yields a subnetwork containing the nodes and edges traversed by the walker.

the network organically come into focus. Consider, for example, the modular network from Fig. 2C. By combining transitions of lengths two, three, or more, humans tend to overweigh the associations within communities and underweigh the transitions between communities (Fig. 3B). This simple observation explains why people are surprised by cross-cluster transitions (42, 43) (Fig. 2D), why sequences in lattice and random networks are more difficult to anticipate (42, 43) (Fig. 2D), and how people detect the boundaries between clusters (39, 41, 57).

More generally, the capacity to learn the large-scale structure of a network enables people to perform many basic cognitive functions, from anticipating nonadjacent dependencies between syllables and words (78, 82) to planning for future events (83, 84) and estimating future rewards (79, 81). Using models similar to that above, researchers have been able to predict the impacts of network structure on human behavior in reinforcement learning tasks (77), pattern detection in random sequences (75, 76), and variations in neural activity (41, 44, 76). Notably, the explained effects span various types of behavioral and neural observations, including reaction times (42, 43, 85), data segmentation (39, 41, 57), task errors (42, 43), randomness detection (86), EEG signals (87), and fMRI recordings (41, 85). Together, these results indicate that people’s ability to detect the mesoscale and global structure of a network emerges not just from their capacity to learn individual edges, but also from their capacity to associate items across spatial, temporal, and topological scales.

The Future of Graph Learning

Past and current advances in graph learning inspire new research questions at the intersection of cognitive science, neuroscience, and network science. Here, we highlight a number of important directions, beginning with possible generalizations of the existing graph learning paradigm before discussing the implications of graph learning for our understanding of the structures and functions of real-world transition networks.

Extending the Graph Learning Paradigm. Most graph learning experiments, including those discussed in Figs. 1 and 2, present each subject with a sequence of stimuli defined by a random walk on a (possibly weighted and directed) transition network (6, 13, 38, 39, 41–47, 57, 78). Equivalently, in the language of stochastic processes, each sequence represents a stationary Markov process (88). Although random walks offer a natural starting point in the study of graph learning, they are also constrained by three main assumptions: 1) that the underlying transition structure remains static over time (stationarity), 2) that future stimuli depend only on the current stimulus (the Markov property), and 3) that the sequence is predetermined without input from the observer. Future graph learning experiments can test the boundaries of these constraints by systematically generalizing the existing graph learning paradigm.

Stationarity. While most graph learning experiments focus on static transition networks, many of the networks that humans encounter in the real world either evolve in time or overlap with other networks in the environment (9, 16, 17, 20, 26). Therefore, rather than simply investigating people’s ability to learn a single network, future experiments should study the capacity for humans to detect the dynamical features of an evolving network (Fig. 4A) or differentiate the distinct features of multiple networks. Early results indicate that, when observing a sequence of stimuli that shifts from one transition structure to another, people’s learned representation of the first network influences their behavior in response to the second network, but that these effects diminish with time (42). This gradual “unlearning” of network structure raises an important question for future research: Rather than investigating how network properties facil-

itate learning—as has been the focus of most graph learning studies—can we determine which properties make a network difficult to forget?

The Markov Property. Thus far, in keeping with the majority of existing graph learning research, we have focused exclusively on sequences in which the next stimulus depends only on the current stimulus; that is, we have focused on sequences that obey the Markov property (88). However, almost all sequences of stimuli or items in the real world involve long-range correlations and dependencies (Fig. 4B). For example, the probability of a word in spoken language depends not just on the previous word, but also on the earlier words in the sentence and the broader context in which the sentence exists (89). Similarly, musical systems often enforce constraints on the length and structure of sequences, thereby inducing long-range dependencies between notes (90). Interestingly, given mounting evidence that people construct long-distance associations (43, 44, 75–78), the resulting internal estimates of transition structures resemble non-Markov processes (43). Therefore, future research could investigate whether the learning of long-distance associations enables people to infer the non-Markov features of sequences in daily life.

Information Seeking. Finally, although many of the sequences that humans observe are prescribed without input from the observer, there are also settings in which people have agency in determining the structure of a sequence. For example, when surfing the Internet (91–94) or following a trail of scientific citations (25), people choose their paths through the underlying hyperlink and citation networks. In this way, people are able to

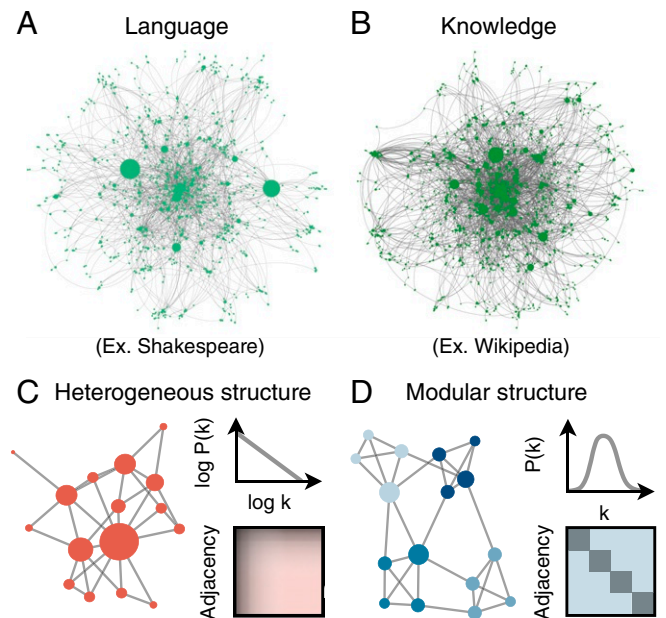


Fig. 5. Real transition networks exhibit hierarchical structure. (A) A language network constructed from the words (nodes) and transitions between them (edges) in the complete works of Shakespeare. (B) A knowledge network of hyperlinks between pages on Wikipedia. (C and D) Many real-world transition networks exhibit hierarchical organization (101), which is characterized by two topological features: (C) Heterogeneous structure, which is often associated with scale-free networks, is typically characterized by a power-law degree distribution and the presence of high-degree hub nodes (55). (D) Modular structure is defined by the presence of clusters of nodes with dense within-cluster connectivity and sparse between-cluster connectivity (21).

seek out information about networks structures rather than simply having the information presented to them (Fig. 4C). Such information seeking has been shown to vary by person (93) and to depend crucially on the topology of the underlying network (91, 92, 94). Moreover, when retrieving information from memory, humans search through their stored networks of associations (95), often performing search strategies that resemble optimal foraging in physical space (96–98). In the context of graph learning, allowing subjects to actively seek information raises a number of compelling questions: Does choosing their path through a transition network enable subjects to more efficiently learn its topology? Or does the ability to seek information lead people to form biased representations of the true transition structure (99, 100)? These questions, combined with the directions described above, highlight some of the exciting extensions of graph learning that will require creative insights and collaborative contributions from cognitive scientists and network scientists alike.

Studying the Structure of Real-World Networks. In addition to shedding light on human behavior and cognition, the study of graph learning also has the promise to offer insights into the structure and function of real-world networks. Indeed, there exists an intimate connection between human cognition and networks: While people rely on networked systems to perform a wide range of tasks, from communicating using language (Fig. 5A) and music to storing and retrieving information through science and the Internet (Fig. 5B), many of these networks have evolved with or were explicitly designed by humans. Therefore, just as humans are adept at learning the structure of networks, one might suspect that some networks are structured to support human learning and cognition.

The perspective that cognition may constrain network structure has recently shed light on the organizational properties of some real-world networks (5, 56), including the small-world structure and power-law degree distributions exhibited by semantic and word co-occurrence networks (16–18) and the scale-free structure of the connections between concepts on Wikipedia (54). Interestingly, many of the networks with which humans interact share two distinct structural features: 1) They are heterogeneous (Fig. 5C), characterized by the presence of hub nodes with unusually high degree (16, 18, 24, 55, 56), and 2) they are modular (Fig. 5D), characterized by the existence of tightly connected clusters (16, 21, 22, 56, 63). Together, heterogeneity and modularity represent the two defining features of hierarchical organization, which has now been observed in a wide array of man-made networks (101, 102). Could it be that the shared structural properties of these networks arise from their common functional purpose: to facilitate human learning and communication?

Graph learning provides quantitative models and experimental tools to begin answering questions such as these (103). For

example, experimental results, such as those discussed in Fig. 2, indicate that modular structure improves people's ability to anticipate transitions (42, 43), and this result has been confirmed numerically using models of the form in Fig. 3 (43). Moreover, the high-degree hubs found in heterogeneous networks have been shown to help people search for information (91, 94). Together, these results demonstrate that graph learning offers a unique and constructive lens through which to study networks in the world around us.

Conclusions and Outlook

Understanding how people learn and represent the complex relationships governing their environment remains one of the greatest open problems in the study of human cognition. On the heels of decades of research in cognitive science and statistical learning investigating how humans detect the local properties of individual items and the connections between them (6, 15, 31–37), conclusive evidence now demonstrates that human behavior, cognition, and neural activity depend critically on the large-scale structure of items and connections (13, 38, 39, 41–44, 57). By casting the items and connections in our environment as nodes and edges in a network, scientists can now explore the impact of network structure on human cognition in a unified and principled framework.

Although the experimental and numerical foundation of the field has been laid, graph learning remains a budding area of research offering a wealth of interdisciplinary opportunities. From cognitive modeling techniques (Fig. 3) and extensions of existing experimental paradigms (Fig. 4) to applications in the study of real-world networks (Fig. 5), graph learning is primed to alter the way we think about human cognition, complex networks, and the myriad ways in which they intersect.

Materials and Methods

The materials and methods discussed in this article are presented and described in the references listed herein.

Data Availability. This article contains no new data.

ACKNOWLEDGMENTS. We thank David Lydon-Staley, Nico Christianson, and Jennifer Stiso for comments on previous versions of this paper. D.S.B. and C.W.L. acknowledge support from the John D. and Catherine T. MacArthur Foundation, the Alfred P. Sloan Foundation, the Institute for Scientific Interchange Foundation, the Paul G. Allen Family Foundation, the Army Research Laboratory (W911NF-10-2-0022), the Army Research Office (Bassett-W911NF-14-1-0679, Grafton-W911NF-16-1-0474, and DCIST-W911NF-17-2-0181), the Office of Naval Research, the National Institute of Mental Health (2-R01-DC-009209-11, R01-MH112847, R01-MH107235, and R21-MH-106799), the National Institute of Child Health and Human Development (1R01HD086888-01), the National Institute of Neurological Disorders and Stroke (R01 NS099348), and the National Science Foundation (BCS-1441502, BCS-1430087, NSF PHY-1554488, and BCS-1631550). The content is solely the responsibility of the authors and does not necessarily represent the official views of any of the funding agencies.

1. E. C. Tolman, Cognitive maps in rats and men. *Psychol. Rev.* **55**, 189–208 (1948).
2. R. G. Golledge, "Human wayfinding and cognitive maps" in *The Colonization of Unfamiliar Landscapes*, M. Rockman, J. Steele, Eds. (Routledge, 2003), pp. 49–54.
3. B. Kosko, Fuzzy cognitive maps. *Int. J. Man Mach. Stud.* **24**, 65–75 (1986).
4. J. Portugali, *The Construction of Cognitive Maps* (Springer Science & Business Media, 1996), vol. 32.
5. A. Baronchelli, R. Ferrer-i Cancho, R. Pastor-Satorras, N. Chater, M. H. Christiansen, Networks in cognitive science. *Trends Cognit. Sci.* **17**, 348–360 (2013).
6. J. R. Saffran, R. N. Aslin, E. L. Newport, Statistical learning by 8-month-old infants. *Science* **274**, 1926–1928 (1996).
7. T. T. Hills, M. M. M. Maouene, A. Sheya, L. Smith, Longitudinal analysis of early semantic networks: Preferential attachment or preferential acquisition? *Psychol. Sci.* **20**, 729–739 (2009).
8. T. Engelthaler, T. T. Hills, Feature biases in early word learning: Network distinctiveness predicts age of acquisition. *Cognit. Sci.* **41**, 120–140 (2017).
9. A. E. Sizemore, E. A. Karuza, C. Giusti, D. S. Bassett, Knowledge gaps in the early growth of semantic feature networks. *Nat. Hum. Behav.* **2**, 682–692 (2018).
10. A. D. Friederici, Neurophysiological markers of early language acquisition: From syllables to sentences. *Trends Cognit. Sci.* **9**, 481–488 (2005).
11. G. Kachergis, C. Yu, R. M. Shiffrin, An associative model of adaptive inference for learning word-referent mappings. *Psychon. Bull. Rev.* **19**, 317–324 (2012).
12. G. Kachergis, C. Yu, R. M. Shiffrin, Actively learning object names across ambiguous situations. *Top. Cognit. Sci.* **5**, 200–213 (2013).
13. S. H. Tompson, A. E. Kahn, E. B. Falk, J. M. Vettel, D. S. Bassett, Individual differences in learning social and nonsocial network structures. *J. Exp. Psychol. Learn. Mem. Cognit.* **45**, 253–271 (2019).
14. W. A. Bousfield, The occurrence of clustering in the recall of randomly arranged associates. *J. Gen. Psychol.* **49**, 229–240 (1953).
15. J. Fiser, R. N. Aslin, Statistical learning of higher-order temporal structure from visual shape sequences. *J. Exp. Psychol. Learn. Mem. Cognit.* **28**, 458–467 (2002).
16. M. Steyvers, J. B. Tenenbaum, The large-scale structure of semantic networks: Statistical analyses and a model of semantic growth. *Cognit. Sci.* **29**, 41–78 (2005).
17. S. N. Dorogovtsev, J. F. Mendes, Language as an evolving word web. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **268**, 2603–2606 (2001).

18. R. F. I. Cancho, R. V. Solé, The small world of human language. *Proc. R. Soc. Lond. B Biol. Sci.* **268**, 2261–2265 (2001).
19. X. F. Liu, K. Tse Chi, M. Small, Complex network structure of musical compositions: Algorithmic generation of appealing music. *Physica A* **389**, 126–132 (2010).
20. A.-L. Barabási *et al.*, Evolution of the social network of scientific collaborations. *Physica A* **311**, 590–614 (2002).
21. M. Girvan, M. E. J. Newman, Community structure in social and biological networks. *Proc. Natl. Acad. Sci. U.S.A.* **99**, 7821–7826 (2002).
22. K. Astrup Eriksen, I. Simonsen, S. Maslov, K. Sneppen, Modularity and extreme edges of the internet. *Phys. Rev. Lett.* **90**, 148701 (2003).
23. R. Albert, H. Jeong, A.-L. Barabási, Internet: Diameter of the world-wide web. *Nature* **401**, 130–131 (1999).
24. M. E. J. Newman, The structure of scientific collaboration networks. *Proc. Natl. Acad. Sci. U.S.A.* **98**, 404–409 (2001).
25. T. Martin, B. Ball, B. Karrer, M. E. J. Newman, Coauthorship and citation patterns in the physical review. *Phys. Rev. E* **88**, 012814 (2013).
26. N. Beckage, L. Smith, T. Hills, Small worlds and semantic network growth in typical and late talkers. *PLoS One* **6**, e19348 (2011).
27. H. Dubossarsky, S. De Deyne, T. T. Hills, Quantifying the structure of free association networks across the life span. *Dev. Psychol.* **53**, 1560–1570 (2017).
28. E. Koehlin, A. Hyafil, Anterior prefrontal function and the limits of human decision-making. *Science* **318**, 594–598 (2007).
29. R. Wilensky, *Planning and Understanding: A Computational Approach to Human Reasoning* (Addison-Wesley, Reading, MA, 1983).
30. M. R. Brent, T. A. Cartwright, Distributional regularity and phonotactic constraints are useful for segmentation. *Cognition* **61**, 93–125 (1996).
31. A. R. Romberg, J. R. Saffran, Statistical learning and language acquisition. *Wiley Interdiscip. Rev. Cogn. Sci.* **1**, 906–914 (2010).
32. R. N. Aslin, E. L. Newport, Statistical learning: From acquiring specific items to forming general rules. *Curr. Dir. Psychol. Sci.* **21**, 170–176 (2012).
33. R. N. Aslin, E. L. Newport, Distributional language learning: Mechanisms and models of category formation. *Lang. Learn.* **64**, 86–105 (2014).
34. A. Schapiro, N. Turk-Browne, “Statistical learning” in *Brain Mapping: An Encyclopedic Reference*, A. W. Toga, Ed. (Elsevier, 2015), pp. 501–506.
35. N. B. Turk-Browne, P. J. Isola, B. J. Scholl, T. A. Treat, Multidimensional visual statistical learning. *J. Exp. Psychol. Learn. Mem. Cogn.* **34**, 399–407 (2008).
36. T. F. Brady, A. Oliva, Statistical learning using real-world scenes: Extracting categorical regularities without conscious intent. *Psychol. Sci.* **19**, 678–685 (2008).
37. D. Baldwin, A. Andersson, J. Saffran, M. Meyer, Segmenting dynamic human action via statistical structure. *Cognition* **106**, 1382–1407 (2008).
38. E. A. Karuza, S. L. Thompson-Schill, D. S. Bassett, Local patterns to global architectures: Influences of network topology on human learning. *Trends Cognit. Sci.* **20**, 629–640 (2016).
39. E. A. Karuza, A. E. Kahn, S. L. Thompson-Schill, D. S. Bassett, Process reveals structure: How a network is traversed mediates expectations about its architecture. *Sci. Rep.* **7**, 12733 (2017).
40. M. E. J. Newman, The structure and function of complex networks. *SIAM Rev.* **45**, 167–256 (2003).
41. A. C. Schapiro, T. T. Rogers, N. I. Cordova, N. B. Turk-Browne, M. M. Botvinick, Neural representations of events arise from temporal community structure. *Nat. Neurosci.* **16**, 486–492 (2013).
42. A. E. Kahn, E. A. Karuza, J. M. Vettel, D. S. Bassett, Network constraints on learnability of probabilistic motor sequences. *Nat. Hum. Behav.* **2**, 936–947 (2018).
43. C. W. Lynn, A. E. Kahn, D. S. Bassett, Abstract representations of events arise from mental errors in learning and memory. *Nat. Commun.*, in press.
44. M. M. Garvert, R. J. Dolan, T. E. J. Behrens, A map of abstract relational knowledge in the human hippocampal–entorhinal cortex. *Elife* **6**, e17086 (2017).
45. A. S. Reber, Implicit learning of artificial grammars. *J. Verb. Learn. Verb. Behav.* **6**, 855–863 (1967).
46. A. Cleeremans, J. L. McClelland, Learning the structure of event sequences. *J. Exp. Psychol. Gen.* **120**, 235–253 (1991).
47. R. L. Gomez, L. A. Gerken, Artificial grammar learning by 1-year-olds leads to specific and abstract knowledge. *Cognition* **70**, 109–135 (1999).
48. J. S. Adelman, G. D. A. Brown, J. F. Quesada, Contextual diversity, not word frequency, determines word-naming and lexical decision times. *Psychol. Sci.* **17**, 814–823 (2006).
49. D. A. Balota, M. J. Cortese, S. D. Sergent-Marshall, D. H. Spieler, M. J. Yap, Visual word recognition of single-syllable words. *J. Exp. Psychol.* **133**, 283–316 (2004).
50. M. T. Carlson, M. Sonderegger, M. Bane, How children explore the phonological network in child-directed speech: A survival analysis of children’s first word productions. *J. Mem. Lang.* **75**, 159–180 (2014).
51. K. Y. Chan, M. S. Vitevitch, The influence of the phonological neighborhood clustering coefficient on spoken word recognition. *J. Exp. Psychol. Hum. Percept. Perform.* **35**, 1934–1949 (2009).
52. R. Goldstein, M. S. Vitevitch, The influence of clustering coefficient on word-learning: How groups of similar sounding words facilitate acquisition. *Front. Psychol.* **5**, 1307 (2014).
53. M. Yates, How the clustering of phonological neighbors affects visual word recognition. *J. Exp. Psychol. Learn. Mem. Cognit.* **39**, 1649–1656 (2013).
54. A. P. Masucci, A. Kalampokis, V. M. Eguiluz, E. Hernández-García, Wikipedia information flow analysis reveals the scale-free architecture of the semantic space. *PLoS One* **6**, e17333 (2011).
55. A.-L. Barabási, R. Albert, Emergence of scaling in random networks. *Science* **286**, 509–512 (1999).
56. J. Borge-Holthoefer, A. Arenas, Semantic networks: Structure and dynamics. *Entropy* **12**, 1264–1302 (2010).
57. E. A. Karuza, A. E. Kahn, D. S. Bassett, Human sensitivity to community structure is robust to topological variation. *Complexity* **2019**, 1–8 (2019).
58. R. Hyman, Stimulus information as a determinant of reaction time. *J. Exp. Psychol.* **45**, 188–196 (1953).
59. G. McCarthy, E. Donchin, A metric for thought: A comparison of p300 latency and reaction time. *Science* **211**, 77–80 (1981).
60. K. I. Forster, S. M. Chambers, Lexical access and naming time. *J. Verb. Learn. Verb. Behav.* **12**, 627–635 (1973).
61. J. R. Anderson, Retrieval of propositional information from long-term memory. *Cognit. Psychol.* **6**, 451–474 (1974).
62. J. R. Anderson, L. M. Reder, The fan effect: New results and new theories. *J. Exp. Psychol.* **128**, 186–197 (1999).
63. A. E. Motter, A. P. S. De Moura, Y.-C. Lai, P. Dasgupta, Topology of the conceptual network of language. *Phys. Rev. E* **65**, 065102 (2002).
64. M. Sigman, G. A. Cecchi, Global organization of the wordnet lexicon. *Proc. Natl. Acad. Sci. U.S.A.* **99**, 1742–1747 (2002).
65. D. J. Watts, S. H. Strogatz, Collective dynamics of ‘small-world’ networks. *Nature* **393**, 440–442 (1998).
66. M. E. J. Newman, Modularity and community structure in networks. *Proc. Natl. Acad. Sci. U.S.A.* **103**, 8577–8582 (2006).
67. J. Moody, Peer influence groups: Identifying dense clusters in large networks. *Soc. Netw.* **23**, 261–283 (2001).
68. M. S. Vitevitch, K. Y. Chan, S. Roodenrys, Complex network structure influences processing in long-term and short-term memory. *J. Mem. Lang.* **67**, 30–44 (2012).
69. K. Y. Chan, M. S. Vitevitch, Network structure influences speech production. *Cognit. Sci.* **34**, 685–697 (2010).
70. J. M. Kleinberg, Navigation in a small world. *Nature* **406**, 845 (2000).
71. M. L. Boas, *Mathematical Methods in the Physical Sciences* (Wiley, 2006).
72. M. W. Howard, M. J. Kahana, A distributed representation of temporal context. *J. Math. Psychol.* **46**, 269–299 (2002).
73. S. Dehaene, F. Meyniel, C. Wacongne, L. Wang, C. Pallier, The neural representation of sequences: From transition probabilities to algebraic patterns and linguistic trees. *Neuron* **88**, 2–19 (2015).
74. F. Meyniel, S. Dehaene, Brain networks for confidence weighting and hierarchical inference during probabilistic learning. *Proc. Natl. Acad. Sci. U.S.A.* **114**, E3859–E3868 (2017).
75. J. Yu Angela, J. D. Cohen, “Sequential effects: Superstition or rational behavior?” in *Advances in Neural Information Processing Systems*, Y. Bengio, Ed. (2009), pp. 1873–1880.
76. F. Meyniel, M. Maheu, S. Dehaene, Human inferences about sequences: A minimal transition probability model. *PLoS Comput. Biol.* **12**, e1005260 (2016).
77. I. Momennejad *et al.*, The successor representation in human reinforcement learning. *Nat. Hum. Behav.* **1**, 680–692 (2017).
78. E. L. Newport, R. N. Aslin, Learning at a distance I. Statistical learning of non-adjacent dependencies. *Cognit. Psychol.* **48**, 127–162 (2004).
79. R. S. Sutton, A. G. Barto, *Introduction to Reinforcement Learning* (MIT Press, Cambridge, MA, 1998), vol. 2.
80. P. Dayan, Improving generalization for temporal difference learning: The successor representation. *Neural Comput.* **5**, 613–624 (1993).
81. S. J. Gershman, C. D. Moore, M. T. Todd, K. A. Norman, P. B. Sederberg, The successor representation and temporal context. *Neural Comput.* **24**, 1553–1568 (2012).
82. G. T. M. Altmann, Y. Kamide, Incremental interpretation at verbs: Restricting the domain of subsequent reference. *Cognition* **73**, 247–264 (1999).
83. C. M. Atance, D. K. O’Neill, Episodic future thinking. *Trends Cognit. Sci.* **5**, 533–539 (2001).
84. D. R. Addis, A. T. Wong, D. L. Schacter, Age-related changes in the episodic simulation of future events. *Psychol. Sci.* **19**, 33–41 (2008).
85. S. A. Huettel, P. B. Mack, G. McCarthy, Perceiving patterns in random series: Dynamic processing of sequence in prefrontal cortex. *Nat. Neurosci.* **5**, 485–490 (2002).
86. R. Falk, C. Konold, Making sense of randomness: Implicit encoding as a basis for judgment. *Psychol. Rev.* **104**, 301–318 (1997).
87. K. C. Squires, C. Wickens, N. K. Squires, E. Donchin, The effect of stimulus sequence on the waveform of the cortical event-related potential. *Science* **193**, 1142–1146 (1976).
88. S. M. Ross *et al.*, *Stochastic Processes* (Wiley, New York, NY, 1996), vol. 2.
89. M. Amit, Y. Shmerler, E. Eisenberg, M. Abraham, N. Shnerb, Language and codification dependence of long-range correlations in texts. *Fractals* **2**, 7–13 (1994).
90. G. R. Jafari, P. Pedram, L. Hedayatifar, Long-range correlation and multifractality in Bach’s inventions pitches. *J. Stat. Mech. Theor. Exp.* **2007**, P04012–P04012 (2007).
91. L. A. Adamic, R. M. Lukose, A. R. Puniyani, B. A. Huberman, Search in power-law networks. *Phys. Rev. E* **64**, 046135 (2001).
92. P. S. Dodds, R. Muhamad, D. J. Watts, An experimental study of search in global social networks. *Science* **301**, 827–829 (2003).
93. V. L. O’Day, R. Jeffries, “Orienting in an information landscape: How information seekers get from here to there” in *Proceedings of the INTERACT’93 and CHI’93 Conference on Human Factors in Computing Systems*, S. Ashlund, K. Mullet, A. Henderson, E. Hollnagel, T. White, Eds. (ACM, New York, NY, 1993), pp. 438–445.

94. R. West, J. Leskovec, "Human wayfinding in information networks" in *Proceedings of the 21st International Conference on World Wide Web*, A. Mille, Ed. (ACM, New York, NY, 2012), pp. 619–628.
95. J. G. Raaijmakers, R. M. Shiffrin, Search of associative memory. *Psychol. Rev.* **88**, 93–134 (1981).
96. T. T. Hills, M. N. Jones, P. M. Todd, Optimal foraging in semantic memory. *Psychol. Rev.* **119**, 431–440 (2012).
97. M. N. Jones, T. T. Hills, P. M. Todd, Hidden processes in structural representations: A reply to Abbott, Austerweil, and Griffiths (2015). *Psychol. Rev.* **122**, 570–574 (2015).
98. J. Avery, M. N. Jones, "Comparing models of semantic fluency: Do humans forage optimally, or walk randomly?" in *Proceedings of the 40th Annual Meeting of the 40th Annual Meeting of the Cognitive Science Society*, T. M. Rodgers, M. Rau, J. Zhu, C. Kalish, Eds. (Curran Associates, Red Hook, NY, 2018).
99. S. Schulz-Hardt, D. Frey, C. Lüthgens, S. Moscovici, Biased information search in group decision making. *J. Pers. Soc. Psychol.* **78**, 655–669 (2000).
100. E. Jonas, S. Schulz-Hardt, D. Frey, N. Thelen, Confirmation bias in sequential information search after preliminary decisions: An expansion of dissonance theoretical research on selective exposure to information. *J. Pers. Soc. Psychol.* **80**, 557–571 (2001).
101. E. Ravasz, A.-L. Barabási, Hierarchical organization in complex networks. *Phys. Rev. E* **67**, 026112 (2003).
102. A. Arenas, A. Fernandez, S. Gomez, Analysis of the structure of complex networks at different resolution levels. *New J. Phys.* **10**, 053039 (2008).
103. C. W. Lynn, L. Papadopoulos, A. E. Kahn, D. S. Bassett, Human information processing in complex networks. *Nat. Phys.*, in press.