

On curiosity: a fundamental aspect of personality, a practice of network growth

Perry Zurn^{1,2} & Danielle S. Bassett^{3,4}

¹Center for Curiosity, University of Pennsylvania, Philadelphia, PA 19104 USA

²Department of Philosophy, American University, Washington, DC 20016 USA

³Department of Bioengineering, University of Pennsylvania, Philadelphia, PA 19104 USA

⁴Department of Electrical & Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104 USA

Abstract

Human personality is reflected in patterns – or *networks* – of behavior, either in thought or action. Curiosity is an oft-treasured component of one’s personality, commonly associated with information-seeking proclivities with distinct neurophysiological correlates. The markers of curiosity can differ substantially across people, suggesting the possibility that personality also determines the architectural style of one’s curiosity. Yet progress in defining those styles, and marking their neurophysiological basis, has been hampered by fairly fundamental difficulties in defining curiosity itself. Here we offer and exercise a definition of the practice of curiosity as knowledge-network building, one particular pattern of thought behavior. To unpack this definition and motivate its utility, we begin with a short primer on network science and describe how the mathematical object of a network can be used to map items and relations that are characteristic of bodies of knowledge. Next, we turn to a discussion of how networks grow, how their growth can be modeled, and how the practice of curiosity can be formalized as a process of network growth. We pay particular attention to how individuals may differ in how they build their knowledge networks, and discuss how the sort, manner, and action of building can be modulated by experience. We discuss how this definition of the practice of curiosity motivates new experiments and theory development at the interdisciplinary intersection of network science, personality neuroscience, education, and curiosity studies. We close with a note on the potential of network science to inform studies of other domains of personality, and the patterns of thought- or action-behavior characteristic thereof.

I. Introduction

A common mantra among practicing psychiatrists is that: “Personality is only behavior, repeated.” While of course a simplification, this notion suggests the possibility that theoretical frameworks to evaluate and summarize behavior might also be relevant for personality studies broadly and personality neuroscience specifically. What are such theoretical frameworks? Behavior, or the expression of the underlying neural system (Hogan, 2015), is classically studied by summarizing the animal’s response to a task in terms of a single variable such as a reaction time or an error rate. Yet, recent advances in imaging approaches, measurement techniques, and storage solutions have supported a marked increase in the acquisition of high-resolution multivariate behavioral data. Hand-curation or machine-learning techniques can distill such data into an interpretable number of basic units of behavior, either for an individual (Egnor & Branson, 2016) or for a collective (Butail, Salerno, Bollt, & Porfiri, 2015). But following such a distillation, it remains difficult to describe the (often probabilistic) sequences of units that comprise an animal’s observable behavior. The difficulty only increases when considering non-observable behaviors such as those that are characteristic of internal thought.

Network science offers a powerful conceptual framework and mathematical formalism for quantitatively describing, modeling, and explaining patterns of behaviors in thought, decision, or action. Whether such patterns arise in the acquisition of motor skills (Wymbs, Bassett, Mucha, Porter, & Grafton, 2012) or linguistic content (Stella, Beckage, & Brede, 2017), or explain retrieval of lexical content from memory (Vitevitch, Chan, & Goldstein, 2014), their complexity is naturally described by the mathematical language of graphs (Bollobás, 2011). A graph is an object in which a system’s fundamental units (nodes) are connected to one another by relations, similarities, or transition probabilities (edges). In this way, skills can be parsimoniously defined as well-learned networks of actions and decisions (Kahn, Karuza, Vettel, & Bassett, n.d.), and thought behaviors can be parsimoniously described as preferred networks of ideas or concepts. Network science also offers a set of tools to describe and model the architectures observed in these networks, thus providing a potentially powerful tool for emerging frontiers in personality neuroscience. As a quintessential example to illustrate this potential, we consider the internal and external behaviors associated with *curiosity*, a particularly revered component of one’s personality, commonly associated with information-seeking proclivities.

Classically, curiosity has been defined, by turns, as a desire for knowledge or a drive for novelty (Aristotle, 1933; Plutarch, 2005). While the former characterization emphasizes the need for satisfaction, such that knowledge would neutralize curiosity, the latter emphasizes the need for stimulation, such that novel information would only intensify curiosity. However useful each definition has been in illuminating the phenomenology of curiosity, both are limited insofar as they characterize curiosity as a motivational state with no reference to individual differences. As such, they support a certain emphasis on incentive over practice and generalization over individuation. Comparatively recent scholarship has begun to analyze curiosity as a function of personality, emphasizing the personal character of our natural inclination to seek information (Gottlieb, Hayhoe, Hikosaka, & Rangel, 2014), whether in our solitary pursuits or our interactions with others. These efforts to formally and systematically determine curiosity’s place as a personality trait have capitalized on experimental approaches and fundamental theories developed in several different scientific disciplines. Neuroscience offers insights into the biological mechanisms of brain processes involved (Kidd & Hayden, 2015). Psychology expands the types of laboratory tests that can probe associated behaviors (Loewenstein, 1994). Neurology and psychiatry address the alterations of

curious thought in disorders of mental health. While each of these domains may define curiosity slightly differently, most assume that information seeking is involved, and that such seeking is driven by motivations that are uniquely calibrated to the person (Oedeyer & Kaplan, 2007) and supported by specific neural processes (Adams, Watson, Pearson, & Platt, 2012).

Defining curiosity around and beyond information seeking requires a distillation of what the latter is, and what it is not. This distillation can sometimes be made easier by first understanding how such seeking might change in different circumstances. Changes in information seeking can occur intrinsically, or can be driven by external factors. Related to perturbative analysis in mathematics (e.g., see Gomez et al., 2013), and to perturbation theory in physics (e.g., see Tsimring & Aranson, 1997), pinpointing the factors that might drive diversity of form is a necessary prerequisite for understanding. Germane to this specific discussion, information seeking takes diverse forms as the brain develops, being driven by complex processes of brain maturation. While children often seek information about novel objects that are “bright, vivid, startling” (Hall & Smith, 1903; James, 1899), adults are more likely to seek information for other reasons, often related to meaning. In addition to the *sort* of information being sought, the *manner* of seeking can differ across people (Amabile, Hill, Hennessey, & Tighe, 1994), over time (Sutin, Beason-Held, Resnick, & Costa, 2009), and in different environments (Aresteh, 1968), spanning from the restless to the abiding. Moreover, the *action* that occurs once the information is found can be quite telling: perhaps bits of information are registered as independent minutiae, or incorporated into one’s existing knowledge by explicit links of analogy, similarity, or causal relation.

Together, the *sort*, *manner*, and *action* make up what we will refer to as the *practice* of curiosity. Colloquially, one might consider browsing in physical or virtual libraries as a commonplace example of the practice of curiosity, where a specific sort of information is sought (something of interest to the browser), in a particular manner (seeking nearby or far-off links between concepts), followed by a distinct action (placing that new information into the browser’s knowledge network by crystalizing the links between the new information and the previously existing information). Yet linking constituents of such a practice, or offering real-world examples of such a practice, is not equivalent to precisely defining it, nor does it provide the exact criteria for quantitative typology, mathematical formulae for modeling, or neuroscientific intuitions for prediction. In this perspective, we ask whether recent advances in interdisciplinary theory related to network science could be useful in formalizing scientific investigations into the practice of curiosity. To address this question, we discuss important recent advances in theoretical constructs, mathematical techniques, and empirical discoveries in network science and neuroscience that support and inform a concrete scientific study of the practice of curiosity. We offer an operational definition of the practice of curiosity as *knowledge network building*. This definition is predicated on the notion that knowledge itself can be represented as a network, and the hypothesis that one can use emerging tools from network science to formally study the manner in which the network grows. The benefits of this operational definition include the ability to connect our intuitions about curiosity with explicit mathematical formulations, to model the functions rather than identify the motivations of curiosity, and thereby to inform new empirical studies of curiosity in humans.

We further explore the implications of this operational definition for understanding curiosity as a critical part of an individual’s personality. If curiosity is the *practice* of knowledge network building, then a curious person engages in this practice more readily, under more conditions, and for longer periods than his or her counterparts. Given, moreover, that personality traits function in tandem, we propose that individual differences uniquely determine the *style* of that practice, one might even

say the architecture of that knowledge network building. Historically, the differential study of curiosity developed only after a general psychological theory had been established (Voss & Keller, 1983). Following Berlyne's influential theory of curiosity as a drive for knowledge acquisition prompted by conceptual conflict and physiological arousal (Berlyne, 1960), researchers correlated it positively with creativity and negatively with anxiety (Leherissey, 1971; Maddi & Berne, 1964; Spielberger & Starr, 1994). Today, scholars continue to explore curiosity's correlation with openness to experience (Kaufman, 2013), appetitive social interactions (Kashdan & Roberts, 2006), and tolerance for uncertainty (Kashdan, Sherman, Yarbro, & Funder, 2013), all while granting that curiosity requires a modicum of anxiety, as a source of optimal stimulus. Importantly for our purposes, Beswick's cognitive process theory of curiosity provides a proto-network approach, defining curiosity as the practice of negotiating novel information and cognitive maps or category systems (Beswick, 1971; Beswick 2017). Highly curious people, in this model, are equally attached to novelty and systematicity, such that they sustain deeper questioning over longer periods of time before resolving novelty into system. In what follows, we similarly explore the individual differences and personality traits that bear on the stylized practice of curiosity as knowledge network building. Broadly speaking, our contribution is thus not a detailed account of an empirical research study, but a philosophical exposition of a new theoretical construct that can synthesize prior observations as well as guide future experiments in the psychology of curiosity.

The remainder of this narrative begins with a discussion of relevant concepts from network science including the elements of a network, mathematical tools to study a network's organization, and more complicated forms of networks, including multilayer networks, that might be important to consider in personality neuroscience broadly, and in curiosity studies specifically. Next, we discuss how knowledge can be represented as a network, and how different sorts of knowledge from concepts to languages have been studied fruitfully from this perspective. We move then to discussing models of network growth as they exist in a few different areas of biology, with a particular focus on models whose formulations contain features relevant for thought itself or the organ of thought (the brain). With the basic concepts and tools in hand, we then offer a detailed exposition of the intersection between the practice of curiosity and models of network growth. We also explore several possible theoretical frames for the intersection of individual differences and knowledge network building. We conclude by outlining current frontiers in theory, computation, and experiment that could deepen our understanding of curiosity (and its neurophysiological basis) as a component of personality, and personality as a determinant of the architectural style of curiosity's knowledge-network building practices. Finally, we provide a brief prologue that offers a speculative discussion of how network science can be used to better understand aspects of personality beyond curiosity.

II. Simple, fundamental concepts from network science

Network science is an emerging discipline that formalizes the study of *complex systems*, which are those composed of many parts whose interactions are heterogeneous enough that they cannot be accurately modeled using mean-field approaches (Newman, 2011). A useful way to represent these systems is in a network, where a part of the system is referred to as a network node (and often visualized as a ball in ball-and-stick diagrams), and where interactions or relations between parts of the system are referred to as network edges (and often visualized as sticks in ball-and-stick diagrams). Networks are naturally mapped onto mathematical objects called graphs (Bollobás, 2002), and can be manipulated and studied straightforwardly by encoding them in *adjacency*

matrices (Newman, 2010). A single adjacency matrix \mathbf{A} is an N -by- N matrix where N is the number of nodes in the graph (or network), and where each ij^{th} element of the matrix gives the strength of the relation between node i and node j . From the structure of such an adjacency matrix, one can learn about a system's organization, make educated guesses about its function, and build simple models of its development, growth, or evolution (Watts & Strogatz, 1998).

When representing a system as a network, the first challenging question facing the curious person is: "What is the right subdivision of this system into parts?" which amounts to asking, "What is the node?" The question is a historical one, being reminiscent of Socrates' penchant for *diairesis*, the practice "of dividing things again by classes, where the natural joints are, and not trying to break any part, after the manner of a bad carver" (Plato 1925, 265e). Answering this question can be difficult because many systems can be subdivided into rather large parts, rather small parts, or any scale in between. It is quite important to surmount this difficulty because from the choice of nodes stems the choice of edges (Butts, 2009): the sort of inter-part relations that can be studied. It is perhaps unsurprising that – in practice – the most useful subdivision of a system into parts depends greatly on the scientific question motivating the study, and whether that question can be isolated to a specific scale of the system's structure or function. In the context of the human brain and human behavior, there certainly exist some cases in which a hypothesis can be fine-tuned to address natural processes occurring at a single scale. However, many other hypotheses cannot be simplified in this way, and must instead be interrogated in a cross-scale manner.

A natural complement to the formulation of a multi-scale hypothesis is the construction of a multiscale network. Take, for example, the primary organ of human curiosity: the human brain can naturally be divided into cortical and subcortical areas whose boundaries can be drawn either by using anatomical information such as cytoarchitecture (Brodmann, 1909), or by using functional information such as patterns of activity (Glasser et al., 2016). Each of these parcels is thought to perform specific computations or produce specific cognitive functions, with subdivisions into larger parcels mapping on to coarser functions and subdivisions into smaller parcels mapping on to finer functions (Betzel & Bassett, 2016). The structural connectivity between parcels can be estimated from diffusion imaging data, which offers information regarding the location and strength of white matter tracts physically linking cortical and subcortical regions (Ghosh & Deriche, 2016). In a complementary analysis, functional connectivity between parcels can be estimated from functional neuroimaging data, which offers information regarding how similar temporal fluctuations in activity are in two parcels (Craddock, Tunngaraza, & Milham, 2015). At smaller spatial scales, structural connections can be defined by synapses between neurons (Kleinfeld et al., 2011); at finer temporal scales, functional connections can be defined as correlations between trains of action potentials (Brody, 1999). The multiscale network of the human brain houses rich information about how complex patterns of thought arise, and how individual differences in neurophysiology can produce individual differences in personality traits such as openness to experience (Beaty, 2016), manners of production such as creativity (Beaty, Benedek, Kaufman, & Silvia, 2015), and behaviors related to knowledge acquisition such as information-seeking (Sherer, Taber-Thomas, & Tranel, 2015).

III. Network representations of knowledge

The complex and multiscale organization of the brain supports our curious human abilities, whereby we secure new knowledge, acquire language(s), and build models of our external world.

What is particularly *apropos* about this fact is the striking conceptual symmetry between the organ of curiosity and the object of curiosity. Indeed, in Webster's 1828 Dictionary of the English language, knowledge is defined as "a clear and certain perception of that which exists, or of truth and fact; *the perception of the connection* and agreement, or disagreement and repugnancy *of our ideas*." Indeed, without adjudicating between correspondence and coherence epistemologies, it is possible to explore individual and collective epistemic networks as belief suites that are produced by both the pressure for beliefs to cohere with one another and the prompt that they correspond to things in the world (Grim, Modell, & Breslin, 2017). In other words, knowledge can be represented as a network of connections (or relations) between ideas, concepts, or bits of information just as the brain can be represented as a network of connections between neural units that process, store, and recall that information.

A simple yet concrete treatment of knowledge as a network can be intuitively considered in the context of language. The units of language span from small phonemes through syllables to words, and each scale at which such network nodes can be defined has a distinct representation in the human cortex (Peeva et al., 2010). At the smallest scale, any two phonemes (two nodes in the language network) can be linked by a network edge weighted by the probability that those two phonemes are found beside one another in a given language. This edge weight can be encoded as the value of the ij^{th} element of the adjacency matrix \mathbf{A} for that language network, and the complete pattern of all pairwise probabilities provides information about the architecture of the language itself. From this information, one can seek to derive rules by which grammars arose in human languages, and one can also define new rules that lead to artificial grammars (Fitch & Friederici, 2012). At a larger scale, any two lexemes (one word or several words, considered as an abstract unit) can be linked if they are phonological neighbors (Luce & Pisoni, 1998), and the network composed of such relations provides important information about how words are retrieved from human memory (Vitevitch, 2008).

Thus, language – our most common formal means of encoding, storing, and transmitting knowledge – can be usefully represented as a multiscale network. But are there natural (yet formal) ways of representing knowledge itself as a network, and is such a representation useful in understanding human curiosity? Initial answers to these questions can be obtained using semantic networks, where nodes represent concepts and edges represent semantic relations between them (Sowa, 1987). Our understanding of the architecture of semantic networks has evolved appreciably over the last several decades. An early study in the 1960s suggested that semantic networks were trees (in the sense of graph theory) (Collins & Quillian, 1969), although a couple studies a decade later suggested that many sets of concepts are not well characterized by such exact hierarchical architecture (Keil, 1979; Slobin, 1973). Capitalizing on expanded computational capabilities and more extensive data collection efforts, studies at the turn of the century replaced the tree hypothesis (Bales & Johnson, 2006) with evidence for two other architectural properties thought to optimize formation and search of semantic memory (Anderson, 2000): (i) small-world organization which can be intuitively described as the presence of strong local clustering alongside a few long-distance connections, and (ii) scale-free organization, indicated by the distribution of node degrees (number of connections emanating from a node) following a power-law (Steyvers & Tenenbaum, 2005).

The exact nature of a semantic network's architecture depends to some degree on the measure of semantic relatedness chosen to define the network's edges. Many possible definitions exist, including those based on estimates of causal relations (Danks, 2014). While including multiple edge

definitions in a single network is possible (using multilayer network approaches [Kievlä et al., 2014]), arguably a simpler place to start is to choose a single, simple measure of semantic relation such as similarity in meaning. Such a choice narrows the object of inquiry from a semantic relatedness network to a semantic similarity network (Harispe, Ranwez, Janaqi, & Montmain, 2015). For example, one could represent single words as nodes, and link two nodes with an edge if they are listed as synonyms in a dictionary; this approach would offer a single semantic similarity network for an entire language, as used by a population. Alternatively, if one is interested in studying individual differences in semantic networks, one could perform a set of behavioral experiments in which human volunteers are asked to list a sequence of words such that temporally contiguous words in the sequence have similar meanings (Kenett, Kenett, Ben-Jacob, & Faust, 2011).

Measuring, quantifying, and understanding the knowledge network of a single human is a necessary first step towards a principled study of how that knowledge network was built, and what role the practice of curiosity played in its construction. Common empirical tools to estimate knowledge networks in a single human include tests requiring verbal responses and tests requiring written responses. Particularly successful examples of the former include free association tasks (Nelson & Zhang, 2000), and tasks in which volunteers are asked to narrate tales either from their imagination or inspired by visual aids; such tasks can also be used to assess alterations in semantic network structure that accompany cognitive deficits associated with neurological disorders or psychiatric disease (Lee et al., 2017; Renz et al., 2003; Drummond et al., 2015; Spitzer 1993). Particularly successful examples of the latter include *post-hoc* analyses of an individual's written work, in any form in which it is accessible. From both verbal and written data, one can also construct a second type of semantic network known as a word co-occurrence network (Lazaridou, Marelli, & Baroni, 2017) in which words (nodes) are connected to one another if they occur less than some number of words away from each other. In principle, this approach could be used to study changes in semantic networks over an author's lifetime, or over the course of a psychiatric disease.

IV. Modeling how networks grow

While understanding the architecture of a human's knowledge network at a single instant in time would be quite useful, that human's personality may be even more associated with how they built that knowledge network up over time. If one wishes to understand the building of a knowledge network, one could consider developing a mathematical model of network growth, and fitting the parameters of such a model to empirically measured data. For example, one might wish to choose a rule by which new nodes in the network are acquired, and one might also wish to choose a rule by which new edges in the network are acquired, either by linking existing nodes, linking new nodes, or linking a new node and an existing node. One might also wish to choose a rule by which nodes or edges are removed (a process related to network *aging*), or in which the strength of edges change (a process related to network *plasticity*). In networks that evolve by both additions and deletions of nodes and edges, it becomes interesting to consider whether there are conservation laws that balance constraints (e.g., energetic or spatial) over time. These considerations are certainly not exhaustive, but form a good set from which to start building a model of network growth.

To gain an intuition for how one might write down a model for the growth of a knowledge network, we first discuss a few canonical benchmark models built on simple topological principles. One of the simplest (and arguably most well-known) network growth model is the preferential attachment

model which was first defined by Price (Price, 1976), and later re-discovered by Barabasi and Albert with slightly different parameter settings (Barabási, 1999). Generally, this network growth model begins with an edge linking node i to node j . Next, a new node is added by placing edges between the new node and m previously existing nodes, where these existing nodes are chosen with a probability related to their degree. The outcome of this process is a network in which a few high degree nodes gain an even greater degree (“the rich get richer”), while the majority of nodes maintain a relatively low degree, as evidenced by a right-skewed degree distribution. The preferential attachment model can be altered by biasing the choice of the m previously existing nodes. For example, the affinity model creates a hierarchical organization by specifying a parameter that takes on a different value at each node; the m existing nodes are then chosen with a probability related to their affinity parameter (Klimm, Bassett, Carlson, & Mucha, 2014).

While the preferential attachment model has proven relevant for understanding the sizes of cities (Simon, 1955), the wealth of rich individuals (Bassett & Bullmore, 2016), the number of citations given to seminal publications (Newman 2002), and the number of links to World Wide Web pages (Barabási, 1999), many real-world and particularly biological systems grow under constraints and pressures that do not allow the formation of very high-degree hubs. In the context of knowledge networks, it is intuitively plausible that such constraints could include limited memory and learning capacity, or spatial encoding constraints in the brain. Indeed, special consideration has been given to adding constraints to network growth models of neural systems, where the probability of two neurons synapsing onto each other depends upon the physical distance between them (Bassett & Bullmore, 2016). One can incorporate this fact in a distance drop-off growth model, where nodes are placed randomly in a Euclidean space, and edges are placed between two nodes with a probability inversely related to their spatial separation (Klimm, Bassett, Carlson, & Mucha, 2014). The outcome of this growth process is a network with a degree distribution that is reminiscent of those observed empirically, and with clear assortativity (Newman, 2002), or the preference for low degree nodes to connect to other nodes of low degree, and high degree nodes to connect to other nodes of high degree. Such models are useful both at the small scale of neurons, and at the large scale of cortical areas (Bullmore & Bassett, 2011).

Constraints on network growth may either directly or indirectly impinge on node-node connectivity. The constraint of physical distance between neurons discussed is an example of a direct constraint. An example of an indirect constraint is the growth (or otherwise temporal variation) of a tissue that either surrounds the network or serves as a substrate for the network. This sort of indirect constraint has recently been modeled in the context of vasculature networks, which are biological distribution systems that transport nutrients across spatial distances to support the health of an organism (Modes, Magnasco, & Katifori, 2016). Here, the growth of the tissue that the vasculature supports is modeled as one dynamical process, and the growth of the vasculature network itself is modeled as a second dynamical process, coupled to the first (Ronellenfitch & Katifori, 2016). Of course, one could also expand the model to include finer-scale biophysical factors as well (Perfahl et al., 2005).

How might these notions of topological wiring rules, physical pressures, and direct *versus* indirect constraints assist us in understanding how knowledge networks grow? Let us first consider this question in the context of semantic networks, where the majority of efforts have focused over the last few years. Steyvers and Tenenbaum (2005) offer a model of semantic network growth whereby nodes (words) are acquired according to a rule that seeks to maximally differentiate nodal patterns of connectivity. Over long time scales, this model produces a network with a scale-free degree

distribution and small-world organization, which they suggest provide good fits to existing empirical data. Hills and colleagues offer a preferential attachment model of word-word association networks to understand how children acquire nouns in the first few years of life (Hills, Maouene, Maouene, Sheya, & Smith, 2009); the model incorporates the frequency with which words are experienced, the diversity of a word's phonological neighbors, and the relation between new words and existing words in the immediate environment. We speculate that such models would be a useful place to start to understand knowledge network building motivated by curiosity. Moreover, although both of these models are suggested to apply to humans generically, it may also be useful to consider variations of these or related models that would fit one individual more than another individual. Such an effort could shed light on individual differences in semantic networks, which are pronounced in empirical studies, and have previously been linked to a participant's personality and creativity (Kenett, Anaki, & Faust, 2014).

V. Informing models of network growth with an individual's practice of curiosity

There are myriad ways in which a knowledge network could grow, and the exact nature of that growth could be passive (a person is exposed to concepts by their caregivers speaking to them) or active (a person purposefully allots 1 hr per day to studying an encyclopedia, and taking vocabulary tests). The previous models of semantic network growth fall in the former category (passive), while here we turn our focus to understanding how the practice of curiosity (active) can grow and reconfigure a single individual's knowledge network. Exactly how such knowledge network reconfiguration should be modeled is an open area of scientific inquiry that will require novel experimental paradigms, as we discuss in more detail later.

Admittedly, a portion of the literature assumes that curiosity is an inborn trait (Litman & Spielberger, 2003), such that one has a natural degree of curiosity just as one has a natural degree of irritability (Stanton & Watson, 2017) and mindfulness (Zhuang et al., 2017). Yet, both irritability and mindfulness vary over time, and mindfulness *training* is becoming an increasingly popular method to alter one's patterns of executive capacities including decision-making (Kirk et al., 2016), memory (Ives-Deliperi, Howells, Stein, Meintjes, & Horn, 2013), and cognitive flexibility (Lee & Orsillo, 2014) by altering brain function (Scheibner, Bogler, Gleich, Haynes, & Bermpohl, 2017). Similarly, curiosity can vary over time (Sternszus, Saroyan, & Steinert, 2017), and can be altered by different environments (Tripathi, Sarkate, Jalgaonkar, & Rege, 2015; Berson & Oreg, 2016), suggesting the possibility of practicing curiosity to change, grow, and enhance one's knowledge network. In his reflections on education, John Dewey suggests that knowledge is a body of learned connections between things. Geographical knowledge entails an understanding of spatial connections, while historical knowledge entails an understanding of human connections (Dewey, 2011). Curiosity, then, is the tendency to make connections between things perceptible. These connections are built in and through experience. It is therefore the educator's task to facilitate experience, encourage curiosity, and thereby enable the growth of knowledge networks in and among their students. But how?

To make this discussion more meaningful, we must first offer an operational definition of the practice of curiosity. Our operational definition is *the performance of mental tasks characteristic of curious thought* (Bassett, forthcoming). Curiosity training could then be said to be a process whereby individuals engage in the practice of certain mental states of curiosity, as well as of certain transitions between states that naturally occur along a direction of questioning. One could practice

engaging different *foci* of curiosity, or performing different types of curious search (Zurn, forthcoming). Perhaps even more fundamentally, one could simply make time to be curious, and to act on curious thoughts. From a network perspective, a person's practice of curiosity amounts to a walk (either random [Pearson, 1905] or biased [Volchenkov & Blanchard, 2011]) along one's knowledge network, choosing which new nodes to acquire and which new edges to add, either as independent units or as subgraphs or motifs (Shen-Orr, Milo, Mangan, & Alon, 2002). One could also model the practice of curiosity as an explicitly physical dynamical process, where one could purposefully leap from their existing knowledge network into an external pool of the knowledge network of the general public (or some canonical "truth" network) in search of a specific expected or non-specific unexpected idea. The way in which a human walks, and the webs that a human seeks to build, are likely to depend in quite a foundational way on our personalities, our prior experiences, and our mental capacities.

Let us now discuss a few important considerations when building a network growth model informed by the practice of curiosity. We must begin by choosing the sort of node for which we search: more likely an idea, or a rather broad concept, than a single word or word form. Next, we must choose the type of edge that we wish to place: perhaps a notion of similarity, or causality, or analogy. It would also be useful to determine if there are characteristics of nodes that would make us more likely to place an edge: for example, perhaps one prefers connecting distant *versus* close-by ideas, or one prefers drawing top-down *versus* bottom-up relations. At a larger scale, we must choose whether and how we will marry the new node and/or new edge with the previously existing network. Are we working towards a dense, highly ordered graph? A sparse, fairly loopy graph? Are the architectural principles of the network consistent across the whole graph? Or do they vary? If they vary, do they vary in some meaningful way with properties of the ideas located in that section of the graph? With Dominic Widdows, we can ask about the distance between ideas, the geometry of the network, and the space in which the network exists (2004), while with Peter Gärdenfors, we can ask: "What is the geometry of curious thought?" (2004), "How does it relate to one's conceptual space?" (2014)

Perhaps we learn how to build knowledge networks from those around us, as we hear them question, we see them search, or we peruse their writings. If so, then the geometries we build may be not altogether unlike the geometries that our mentors and compatriots build. This symmetry of geometries is likely to be the case if knowledge network building is learned in the same way in which we learn the statistics of our environment (Rebuschat & Williams, 2012), a process supported by neural computations that encode either pairwise or higher-order relationships between concepts (Karuza, Thompson-Schill, & Bassett, 2016). By learning the practice of curiosity from other humans (Engel, 2015), we learn which sorts of ideas others seek, and how they link those ideas together into larger and larger webs. The speed and success with which we learn the practice of curiosity from others could depend on the degree of reinforcement provided following our decisions and actions. Reinforcement learning, in which a learner is given feedback about the accuracy of their responses (Schultz, 2015), could serve to strengthen curious behavior as it is validated by another person.

To the degree that the practice of curiosity can be learned from others, it would be important to understand how current educational practice might already support that learning. Further, it would be important to determine whether and how we could alter educational practice to better imbue youth with a healthy practice of curiosity, according to their proclivities and appropriate to their stage of life. Are there optimal ways of transmitting information about what sorts of new network

nodes to search for, what sorts of edges to use to link them, what sorts of architectures to fill out? Is the traditional lecture format conducive to such transmission? Would mentorship or apprenticeship approaches be more effective? Should the tutelage be implicit (merely viewing the practice of curiosity in others), or explicit (being told that one is learning how to practice curiosity), or both? Answering these questions will require additional work at the interdisciplinary intersection of network science, personality neuroscience, education, and curiosity studies.

VI. Bridging from curious knowledge network building to underlying neurophysiology

What are the neural processes that support curious knowledge network building? Answering this question requires that we first consider neural representations of knowledge networks. Most early work in mapping representations of concepts or objects focused on single items, such as comparing and contrasting the representation of a house vs. that of a face (O'Craven, Downing, & Kanwisher, 1999). This work motivated broader questions about how the representations between different objects or concepts are related to one another, and how that similarity might impact memory and by extension the growth of one's knowledge space (see Xue et al., 2010 for an early empirical study and Haxby, Connolly, & Guntupalli, 2014 for a recent review). In a pioneering study, Behrens and colleagues recently demonstrated that the brain can organize concepts into a 2-dimensional mental map, allowing conceptual relationships to be navigated in a hexagonal grid-like network code (Constantinescu, O'Reilly, & Behrens, 2016), similar to the manner in which humans navigate space (Hafting, Fyhn, Molden, Moser, & Moser, 2005). Exactly how that organization takes place – or how we learn to represent concepts in the brain and to encode their network of relationships – is an active area of inquiry. One posited mechanism for such learning is temporal contingencies (Karuza, Thompson-Schill, & Bassett, 2016). Indeed, recent evidence suggests that humans can infer the community structure present in networks of visual stimuli, when those stimuli are shown to the participant in a continuous stream defined by a walk on the network (Schapiro, Rogers, Cordova, Turk-Browne, & Botvinick, 2013). Follow-up studies demonstrated that learning could be modulated by the type of walk taken through the network (Karuza, Kahn, Thompson-Schill, & Bassett, 2017), and tracked by the representations, activity dynamics, and connectivity of the hippocampus (Schapiro, Turk-Browne, Norman, & Botvinick, 2016).

The fact that the learning and appropriation of network structure in stimuli can be modulated by the temporal sequence of stimuli supports our hypothesis that one's practice of curiosity (or the manner in which one samples knowledge space) will impact the type of knowledge network built. Yet, an important open question lies in what neurophysiological mechanisms might drive individual differences in these processes. In a comprehensive and thought-provoking recent review, Kidd and Hayden summarize recent work probing the biological function, mechanisms, and neural underpinning of curiosity (2015). This body of work links curiosity with behavioral characteristics that differ across individuals, including our preferences for novelty vs. familiarity (Kang et al., 2009), for exploring our environment in search of new information vs. exploiting existing information (Daw, O'Doherty, Dayan, Seymour, & Dolan, 2006), and for the temporal resolution of uncertainty (Blanchard, Hayden, & Bromberg-Martin, 2015). While the neural mechanisms of curious behaviors are not fully understood, one particularly interesting line of evidence shows that self-reported curiosity tracks the activation of the caudate nucleus and the inferior frontal gyrus, both key players in the human reward circuit, whose function is heavily dependent on the neurotransmitter dopamine (Kang et al., 2009). Interestingly, the nucleus accumbens – also consistently implicated in reward circuitry (Knutson, Adams, Fong, & Hommer, 2001) – was not activated. When an answer to the question that the participants were curious about was revealed,

areas involved in learning and memory (such as parahippocampal gyrus and hippocampus) were activated, indicating the tight link between curiosity and learning (Kang et al., 2009).

These separate lines of inquiry motivate future experiments tracking neural representations of concepts as human participants seek information about those concepts and interconnections between concepts, as well as experiments addressing the question of how individual differences in (i) the acquisition and change in neural representations, (ii) preferences for novelty, exploration, and uncertainty, and (iii) reward circuit activity might relate to one another.

VII. Current Frontiers

In the previous sections, we have reviewed a range of empirical studies across disciplinary boundaries that inform our formal framing of the practice of curiosity in terms of the mathematics of network science. However, we note that direct empirical studies that capitalize on this framework have yet to be undertaken. It is natural to begin by measuring the characteristics of human objects of curiosity. Are there commonalities between them? Are we more likely to search for an idea or a causal relation? How well can we articulate what we are searching for, verbally, pictorially, or in written form? What sorts of relations between these objects are we most likely to seek? And how does the topology of our network change as we practice curiosity? To what degree does curiosity add to versus reconfigure our knowledge networks? How are these knowledge networks, or their architectural properties, represented in the brain?

One fairly straightforward way to begin answering these questions is to conduct laboratory experiments with living human volunteers in which we measure evolving knowledge networks. Such studies could evaluate changes in knowledge networks during explicit curiosity training or during intrinsic information seeking with little or no external constraints. These data could be used to test the hypothesis that individual differences in the practice of curiosity are reflected in individual differences in the knowledge networks that are grown. Further, one could test the hypothesis that the manner in which knowledge networks are grown is correlated with more commonly studied aspects of one's personality. During such studies, one could also consider mapping the neural correlates of different sorts of walks along knowledge networks, both in their baseline form, and as they evolve over time. One could use these data to test the hypothesis that circuitry supporting reward processing and information-seeking behavior differ across individuals in a manner that maps onto individual differences in their practice of curiosity. A complementary alternative is to measure the evolution of knowledge networks in the written work of moderately prolific authors. While laboratory experiments could provide controlled environments in which to modulate curiosity training and knowledge network growth, and to link this information to underlying neurophysiology, they will likely be constrained to short time scales. By contrast, studies of authors (particularly those who are deceased) could provide more ecologically valid accounts of knowledge network growth over longer time scales.

Experimental studies could be usefully complemented by theoretical studies developing simple mathematical models of the observed phenomena. The models currently developed to characterize the acquisition of nouns in young children may contain rules that are also relevant for the growth of knowledge networks in adult humans actively practicing curiosity. Yet, it is also likely that the network evolution characteristic of the practice of curiosity will require additional or altogether different growth rules mapping the purposeful reaching for distant ideas. In any case, such models

could be parameterized directly from the empirical data. It would be particularly useful to determine whether best fitting parameter values differ for different individuals, at different stages of their life, with different brain architectures, and across individuals experiencing different levels of mental health.

VIII. Summary

In summary, this perspective seeks to develop and offer a formal framework for the study of curiosity as embedded in the joint disciplines of personality neuroscience and network science. We provide a short primer on network science and describe how the mathematical object of a graph or network can be used to map the items and their relations that are characteristic of bodies of knowledge, real and artificial grammars, and language. We discuss how networks can grow, and a few simple models that have been developed to better understand growth in various networks, including semantic networks. We offer an operational definition of the practice of curiosity, and discuss how that practice can be formalized as a process of network growth. We place particular emphasis on individual differences in knowledge network growth, informed by various aspects of personality and by our previous and current educational experiences. We outline important current frontiers in empirical measurement and theoretical model validation with the aim of developing concrete training regimens and educational practices in the future to enhance the practice of curiosity in youth.

IX. Epilogue: Beyond curiosity – The general relevance of network science for personality neuroscience

We began this article by noting that the behaviors reflecting our personalities, and the neural drivers thereof, are fundamentally multivariate and characterized by complex temporal patterns of transitions between behavioral units. Moreover, we posited that network science offers a powerful conceptual framework and mathematical formalism for quantitatively describing, modeling, and explaining such patterns of behaviors, whether present in the domains of thought, decision, or action. The majority of our exposition focused on one particular domain of personality: curiosity, composed of both internal and external behaviors related to the acquisition of networks of ideas or concepts. Yet, the framework and formalism of network science are relevant across other domains of personality neuroscience more broadly.

Perhaps the most direct extension of the ideas we have described here lies in considering other time-varying behaviors of single humans. For example, one could use network science to understand the manner in which cognitive, interpersonal, or affective states relate to one another, similarly to the manner in which knowledge states relate to another. One could also extend these tools beyond a single individual and consider a group of individuals: what is the architecture of the network in which individuals (nodes) are connected by similarities in their personality traits (edges)? Does a person's placement in that network relate to real-world outcomes? Finally, one could consider expanding beyond a single network to understand the relationships between networks. In probing cognitive domains of personality, one can ask how one's internal knowledge network is instantiated in one's physical brain network. In probing sociocultural domains of personality, one can ask how one's brain network relates to one's placement in their social network (Falk & Bassett, 2017). In the future, we envision that such studies of individual, group, and multi-scale network architectures across various domains of personality will be increasingly important for understanding the structure and neurophysiological drivers of the human condition.

X. Acknowledgments

The authors gratefully acknowledge helpful conversations with the founder and staff of the Center for Curiosity, University of Pennsylvania, as well as with the participants of the recent symposium “Curiosity Across the Disciplines Symposium” (Dec. 9, 2016), particularly Susan Engel and Barbara Benedict. DSB acknowledges support from the Alfred P. Sloan Foundation, the John D. and Catherine T. MacArthur Foundation, and NSF CAREER PHY-1554488. The content is solely the responsibility of the authors and does not necessarily represent the official views of any of the funding agencies.

References

- Adams, G. K., Watson, K. K., Pearson, J., Platt, M. L. (2012). Neuroethology of decision-making. *Current Opinion in Neurobiology*, 22(6), 982-9. doi: 10.1016/j.conb.2012.07.009.
- Amabile, T. M., Hill, K. G., Hennessey, B. A., Tighe, E. M. (1994). The work preference inventory: assessing intrinsic and extrinsic motivational orientations. *Journal of Personality and Social Psychology*, 66(5), 950-67. doi=10.1.1.596.7369.
- Anderson, J. R. (2000). *Learning and memory: An integrated approach* (2nd ed.). New York: Wiley.
- Arasteh, J. (1968). Creativity and related processes in the young child: a review of the literature. *Journal of Genetic Psychology*, 112(1st Half), 77-108.
- Aristotle (1993). *Metaphysics*. Cambridge: Loeb Classical Library.
- Bales, M. E., Johnson, S. B. (2006). Graph theoretic modeling of large-scale semantic networks. *Journal of Biomedical Informatics*, 39(4), 451-64. doi:10.1016/j.jbi.2005.10.007.
- Barabási, A.-L. & Albert, R. (1999). Emergence of scaling in random networks. *Science*, 286(5439), 509–512. doi: 10.1126/science.286.5439.509.
- Bassett, D. S. & Bullmore, E. T. (2016). Small-world brain networks revisited. *Neuroscientist*, pii: 1073858416667720, Epub ahead of print. doi: 10.1177/1073858416667720.
- Bassett, D. S. (forthcoming). A network science of the practice of curiosity. In *Curiosity Studies: Toward a New Ecology of Knowledge*, eds. P. Zurn & A. Shankar, Minneapolis: University of Minnesota Press.
- Beaty, R. E., Benedek, M., Kaufman, S. B., Silvia, P. J. (2015). Default and executive network coupling supports creative idea production. *Scientific Reports*, 5, 10964. doi: 10.1038/srep10964.
- Beaty, R. E., Kaufman, S. B., Benedek, M., Jung, R. E., Kenett, Y. N., Jauk, E., Neubauer, A. C., Silvia, P. J. (2016). Personality and complex brain networks: The role of openness to experience in default network efficiency. *Human Brain Mapping*, 37(2), 773-9. doi: 10.1002/hbm.23065.
- Berlyne, D. (1960). *Conflict, Arousal, & Curiosity*. New York: McGraw-Hill Publishers.
- Berson, Y. & Oreg, S. (2016). The role of school principals in shaping children's values. *Psychological Science*, 27(12), 1539-1549. doi: 10.1177/0956797616670147.
- Beswick, D. G. (1971). Reexamination of two learning style studies in the light of the cognitive process theory of curiosity. *Journal of Educational Psychology* 62, 456-462. doi:10.1037/h0031817.
- Beswick, D. G. (2017). *Cognitive motivation: from curiosity to identity, purpose, and meaning*. Cambridge: Cambridge University Press. doi:10.1017/9781316822920.
- Betz, R. F. & Bassett, D. S. (2016). Multi-scale brain networks. *Neuroimage*, S1053-8119(16), 30615-2. doi: 10.1016/j.neuroimage.2016.11.006.
- Blanchard, T. C., Hayden, B. Y., Bromberg-Martin, E. S. (2015). Orbitofrontal cortex uses distinct codes for different choice attributes in decisions motivated by curiosity. *Neuron*, 85, 602–614. doi:10.1016/j.neuron.2014.12.050.
- Bollobás, B. (2002). *Modern graph theory*. New York: Springer. 10.1007/978-1-4612-0619-4.
- Bollobás, B. (2011). *Random graphs*. Cambridge: Cambridge University Press. doi:10.1017/CBO9780511814068.
- Brodman, K. (1909). "Vergleichende lokalisationslehre der grosshirnrinde." Leipzig: Johann Ambrosius Barth. doi: 10.1007/b138298.
- Brody, C. D. (1999). Correlations without synchrony. *Neural Computation*, 11(7), 1537-51. doi: 10.1162/089976699300016133.
- Bullmore E. T. & Bassett, D. S. (2011). Brain graphs: graphical models of the human brain connectome. *Annual Review of Clinical Psychology*, 7, 113-40. doi:10.1146/annurev-clinpsy-040510-143934.

- Butail, S., Salerno, P., Bollt, E. M., Porfiri, M. (2015). Classification of collective behavior: a comparison of tracking and machine learning methods to study the effect of ambient light on fish shoaling. *Behavioral Research Methods*, 47(4), 1020-31. doi:10.3758/s13428-014-0519-2.
- Butts, C. T. (2009). Revisiting the foundations of network analysis. *Science*, 24(325), 414-6. doi: 10.1126/science.1171022.
- Collins, A. M. & Quillian, M. R. (1969). Retrieval time from semantic memory. *Journal of Verbal Learning and Verbal Behavior*, 8, 240–248. Doi:10.1016/S0022-5371(69)80069-1
- Constantinescu, A. O., O'Reilly, J. X., Behrens, T. E. J. (2016). Organizing conceptual knowledge in humans with a gridlike code. *Science*, 352(6292), 1464-1468. doi: 10.1126/science.aaf0941.
- Craddock, R. C., Tunjaraza, R. L., Milham, M. P. (2015). Connectomics and new approaches for analyzing human brain functional connectivity. *Gigascience*. 25(4):13. doi: 10.1186/s13742-015-0045-x.
- Danks, D. (2014). *Unifying the mind: cognitive representations as graphical models*. Boston: MIT Press.
- Daw, N. D., O'Doherty, J. P., Dayan, P., Seymour, B., Dolan, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature*, 441, 876–879. doi: 10.1038/nature04766.
- Dewey, J. (2011). *Democracy and education*. New York: Simon and Brown. doi:10.1086/591813.
- Drummond, C., Coutinho, G., Fonseca, R. P., Assunção, N., Teldeschi, A., de Oliveira-Souza, R., Moll, J., Tovar-Moll, F., Mattos, P. (2015). Deficits in narrative discourse elicited by visual stimuli are already present in patients with mild cognitive impairment. *Frontiers in Aging Neuroscience*, 7, 96. doi: 10.3389/fnagi.2015.00096
- Egnor, S. E. & Branson, K. (2016). Computational analysis of behavior. *Annual Review of Neuroscience*, 8 (39), 217-36. doi:10.1146/annurev-neuro-070815-013845.
- Engel, S. (2015). *The hungry mind: the origins of curiosity in childhood*. Cambridge: Harvard University Press.
- Falk, E. B. & Bassett, D. S. (2017). Brain and social networks: fundamental building blocks of human experience. *Trends in Cognitive Science*, 21(9), 674-690. doi: 10.1016/j.tics.2017.06.009.
- Fitch, W. T. & Friederici, A.D. (2012). Artificial grammar learning meets formal language theory: an overview. *Philosophical Transactions: The Royal Society of London, Biological Sciences*, 367(1598), 1933-55. doi: 10.1098/rstb.2012.0103.
- Gärdenfors, P. (2004). *Conceptual spaces: the geometry of thought*. Boston: MIT Press.
- Gärdenfors, P. (2014). *The geometry of meaning: semantics based on conceptual spaces*. Boston: MIT Press.
- Glasser, M. F., Coalson, T. S., Robinson, E. C., Hacker, C. D., Harwell, J., Yacoub, E., Ugurbil, K., Andersson, J., Beckmann, C. F., Jenkinson, M., Smith, S. M., Van Essen, D. C. (2016). A multi-modal parcellation of human cerebral cortex. *Nature*, 536(7615), 171-8. doi: 10.1038/nature18933.
- Ghosh, A. & Deriche, R. (2016). A survey of current trends in diffusion MRI for structural brain connectivity. *Journal Neural Engineering*, 13(1), 011001. doi: 10.1088/1741-2560/13/1/011001.
- Gómez, S., Díaz-Guilera, A., Gómez-Gardeñes, J., Pérez-Vicente, C. J., Moreno, Y., Arenas, A. (2013). Diffusion dynamics on multiplex networks. *Physical Review Letters*, 110, 028701. doi: 10.1103/PhysRevLett.110.028701
- Gottlieb, J., Hayhoe, M., Hikosaka, O., Rangel, A. (2014). Attention, reward, and information seeking. *Journal of Neuroscience*, 34(46), 15497-504. doi: 10.1523/JNEUROSCI.3270-14.2014.

- Grim, P., Modell, A., Breslin, N. (2017). Coherence and correspondence in the network dynamics of belief suites. *Episteme: A Journal of Individual and Social Epistemology* 14(2), 233-253. doi: 10.1017/epi.2016.7
- Hafting, T., Fyhn, M., Molden, S., Moser, M. B., Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, 436(7052), 801-6. doi:10.1038/nature03721.
- Hall, G. S. & Smith, T. L. (1903). Curiosity and interest. *Pedagogical Seminary*, 10, 315-358. doi:10.1080/08919402.1903.10532722.
- Harispe S., Ranwez, S., Janaqi, S., Montmain, J. (2015). Semantic similarity from natural language and ontology analysis. *Synthesis Lectures on Human Language Technologies*, 8(1), 1-254. ; doi:10.2200/S00639ED1V01Y201504HLT027.
- Haxby, J. V., Connolly, A. C., Guntupalli, J. S. (2014). Decoding neural representational spaces using multivariate pattern analysis. *Annual Review of Neuroscience*, 37, 435-56. doi: 10.1146/annurev-neuro-062012-170325.
- Hills, T. T., Maouene, M., Maouene, J., Sheya, A., Smith, L. (2009). Longitudinal analysis of early semantic networks: preferential attachment or preferential acquisition? *Psychological Science*, 20(6), 729-39. doi: 10.1111/j.1467-9280.2009.02365.x.
- Hogan, J. A. (2015). A framework for the study of behavior. *Behavioral Processes*, 117, 105-13. doi: 10.1016/j.beproc.2014.05.003
- Holger Perfahl, H., Hughes, B. D., Alaracon, R., Maini, P. K., Lloyd, M. C., Reuss, M., Byrne, H. M. (2017). 3D hybrid modeling of vascular network formation. *Journal of Theoretical Biology*. 414(, 254-268. doi: 10.1016/j.jtbi.2016.11.013.
- Ives-Deliperi, V. L., Howells, F., Stein, D. J., Meintjes, E. M., Horn, N. (2013). The effects of mindfulness-based cognitive therapy in patients with bipolar disorder: a controlled functional MRI investigation. *Journal of Affective Disorders*, 150(3), 1152-7. doi: 10.1016/j.jad.2013.05.074.
- James, W. (1899). *Talks to teachers on psychology: and to students on some of life's ideals*. New York: Henry Holt & Company. doi: 10.1037/10814-000.
- Kahn, A. E., Karuza, E. A., Vettel, J. M., Bassett, D. S. (N.D.). Network constraints on learnability of probabilistic motor sequences. arXiv:1709.03000.
- Kang, M. J., Hsu, M., Krajbich, I. M., Loewenstein, G., McClure, S. M., Wang, J. T. Y., Camerer, C. F. (2009). The wick in the candle of learning: epistemic curiosity activates reward circuitry and enhances memory. *Psychological Science*, 20, 963-973. doi: 10.1111/j.1467-9280.2009.02402.x.
- Karuza, E. A., Thompson-Schill, S. L., Bassett, D. S. (2016). Local patterns to global architectures: Influences of network topology on human learning. *Trends in Cognitive Science*, 20(8), 629-40. doi: 10.1016/j.tics.2016.06.003.
- Karuza, E. A., Kahn, A. E., Thompson-Schill, S. L., Bassett, D. S. (2017). Process reveals structure: How a network is traversed mediates expectations about its architecture. *Scientific Reports*, 7(1), 12733. doi: 10.1038/s41598-017-12876-5.
- Kashdan, T. B. & Roberts, J. E. (2006). Affective outcomes in superficial and intimate interactions: Roles of social anxiety and curiosity. *Journal of Research in Personality*, 40, 140-167. Doi: 10.1016/j.jrp.2004.10.005.
- Kashdan, T. B., Sherman, R. A., Yarbro, J., & Funder, D. C. (2013). How are curious people viewed and how do they behave in social situations? From the perspectives of self, friends, parents, and unacquainted observers. *Journal of Personality*, 81(2), 142-154. doi: 10.1111/j.1467-6494.2012.00796.x.

- Kaufman, S. B. (2013). Opening up openness to experience: A Four-factor model and relations to creative achievement in the arts and sciences. *Journal of Creative Behavior*, 47(4), 233-255. doi: 10.1002/jocb.33.
- Keil, F. C. (1979). *Semantic and conceptual development: an ontological perspective*. Cambridge, MA: Harvard University Press.
- Kenett, Y. N., Kenett, D. Y., Ben-Jacob, E., Faust, M. (2011). Global and local features of semantic networks: evidence from the Hebrew mental lexicon. *PLoS One*, 6(8), e23912. doi: 10.1371/journal.pone.0023912.
- Kenett, Y. N., Anaki, D., Faust, M. (2014). Investigating the structure of semantic networks in low and high creative persons. *Frontiers in Human Neuroscience*, 8, 407. doi: 10.3389/fnhum.2014.00407.
- Kidd, C. & Hayden, B. Y. (2015). The psychology and neuroscience of curiosity. *Neuron*, 88(3), 449-60. doi: 10.1016/j.neuron.2015.09.010.
- Kirk, U., Gu, X., Sharp, C., Hula, A., Fonagy, P., Montague, P. R. (2016). Mindfulness training increases cooperative decision making in economic exchanges: Evidence from fMRI. *Neuroimage*, 138, 274-83. doi: 10.1016/j.neuroimage.2016.05.075.
- Kivelä, M., Arenas, A., Barthelemy, M., Gleeson, J. P., Moreno, Y., Porter, M. A. (2014). Multilayer networks. *Journal of Complex Networks*, 2(3), 203-271. doi: 10.1093/comnet/cnu016.
- Kleinfeld, D., Bharioke, A., Blinder, P., Bock, D. D., Briggman, K. L., Chklovskii, D. B., Denk, W., Helmstaedter, M., Kaufhold, J. P., Lee, W. C., Meyer, H. S., Micheva, K. D., Oberlaender, M., Prohaska, S., Reid, R. C., Smith S. J., Takemura, S., Tsai, P. S., Sakmann B. (2011). Large-scale automated histology in the pursuit of connectomes. *Journal of Neuroscience* 31(45), 16125-38. doi: 10.1523/JNEUROSCI.4077-11.2011.
- Klimm, F., Bassett, D. S., Carlson, J. M., Mucha, P. J. (2014). Resolving structural variability in network models and the brain. *PLoS Computational Biology*, 10(3), e1003491. doi: 10.1371/journal.pcbi.1003491.
- Knutson, B., Adams, C. M., Fong, G. W., Hommer, D. (2001). Anticipation of increasing monetary reward selectively recruits nucleus accumbens. *Journal of Neuroscience*, 21, RC159.
- Lazaridou, A., Marelli, M., Baroni, M. (2017). Multimodal word meaning induction from minimal exposure to natural text. *Cognitive Science*, Epub ahead of print. doi: 10.1111/cogs.12481.
- Lee, J. K. & Orsillo, S. M. (2014). Investigating cognitive flexibility as a potential mechanism of mindfulness in generalized anxiety disorder. *Journal of Behavior Therapy and Experimental Psychiatry*, 45(1), 208-16. doi: 10.1016/j.jbtep.2013.10.008.
- Lee, M., Martin, G. E., Hogan, A., Hano, D., Gordon, P. C., Losh, M. (2017). What's the story? A computational analysis of narrative competence in autism. *Autism*, 1362361316677957. doi: 10.1177/1362361316677957.
- Leherissey, B. (1971). The development of a measure of state epistemic curiosity. Technical Memorandum 34, Florida State University.
- Litman, J. A. & Spielberger, C. D. (2003). Measuring epistemic curiosity and its diversive and specific components. *Journal of Personality Assessment*, 80(1), 75-86. doi: 10.1207/S15327752JPA8001_16.
- Loewenstein, G. (1994). The psychology of curiosity: a review and reinterpretation. *Psychological Bulletin*, 116, 75-98. doi: 10.1037/0033-2909.116.1.75.
- Luce, P. A., Pisoni, D. B. (1998). Recognizing spoken words: the neighborhood activation model. *Ear Hear*, 19(1),1-36.
- Maddi, S. R. & Berne, N. (1964). Novelty of productions and the desire for novelty as active and passive forms of the need for variety. *Journal of Personality*, 32, 270-277.

- Modes, C. D., Magnasco, M. O., Katifori, E. (2016). Extracting hidden hierarchies in 3D distribution networks. *Physical Review X*, 6, 031009. doi: 10.1103/PhysRevX.6.031009.
- Nelson, D. L. & Zhang, N. (2000). The ties that bind what is known to the recall of what is new. *Psychonomic Bulletin & Review*, 7(4), 604-17. doi: 10.3758/BF03212998.
- Newman, M. E. (2002). Assortative mixing in networks. *Physical Review Letters*, 89(20), 208701. doi: 10.1103/PhysRevLett.89.208701.
- Newman, M. E. J. (2010). *Networks: an introduction*. Oxford: Oxford University Press. doi:10.1093/acprof:oso/9780199206650.001.0001.
- Newman, M. E. J. (2011). Complex systems: a survey. *American Journal of Physics*, 79, 800-810. doi: 10.1119/1.3590372.
- O'Craven, K. M., Downing, P. E., Kanwisher, N. (1999). fMRI evidence for objects as the units of attentional selection. *Nature*, 401(6753), 584-7. doi: 10.1038/44134.
- Oedeyer, P.-Y. & Kaplan, F. (2007). What is intrinsic motivation? A typology of computational approaches. *Frontiers in Neurobotics*, 1, 6. doi: 10.3389/neuro.12.006.2007.
- Pearson, K. (1905). The problem of the random walk. *Nature*, 72(1865), 294. doi: 10.1038/072294b0
- Peeva, M. G., Guenther, F. H., Tourville, J. A., Nieto-Castanon, A., Anton J. L., Nazarian, B., Alario, F.X. (2010). Distinct representations of phonemes, syllables, and supra-syllabic sequences in the speech production network. *Neuroimage*, 50(2), 626-38. doi: 10.1016/j.neuroimage.2009.12.065.
- Plato (1925). Phaedrus. In *Euthyphro, Apology, Crito, Phaedo, Phaedrus*. Cambridge: Harvard University Press.
- Plutarch (2005). On being a busybody. *Moralia VI*. Cambridge: Loeb Classical Library, 473-517.
- Price, D. J. de S. (1976). A general theory of bibliometric and other cumulative advantage processes. *Journal of American Society for Information Science*, 27 (5), 292-306. doi: 10.1002/asi.4630270505.
- Rebuschat, P. & Williams, J. N. (2012). *Statistical learning and language acquisition*. Berlin: Walter de Gruyter.
- Renz, K., Lorch, E. P., Milich, R., Lemberger, C., Bodner, A., Welsh, R. (2003). On-line story representation in boys with attention deficit hyperactivity disorder. *Journal of Abnormal Child Psychology*, 31(1), 93-104. doi: 10.1023/A:1021777417160.
- Ronellenfitsch, H. & Katifori, E. (2016). Global optimization, local adaptation, and the role of growth in distribution networks. *Physical Review Letters*, 117(13), 138301. doi: 10.1103/PhysRevLett.117.138301.
- Scherer, A. M., Taber-Thomas, B. C., Tranel, D. (2015). A neuropsychological investigation of decisional certainty. *Neuropsychologia*, 70, 206-13. doi: 10.1016/j.neuropsychologia.2015.02.036.
- Schultz, W. (2015). Neuronal reward and decision signals: from theories to data. *Physiological Review*, 95(3), 853-951. doi: 10.1152/physrev.00023.2014.
- Shen-Orr, S. S., Milo, R., Mangan, S., Alon, U. (2002). Network motifs in the transcriptional regulation network of *Escherichia coli*. *Nature Genetics*, 31(1): 64-8. doi: 10.1038/ng881.
- Schapiro, A. C., Rogers, T. T., Cordova, N. I., Turk-Browne, N. B., Botvinick, M. M. (2013). Neural representations of events arise from temporal community structure. *Nature Neuroscience*, 16(4), 486-92. doi: 10.1038/nn.3331.
- Schapiro, A. C., Turk-Browne, N. B., Norman, K. A., Botvinick, M. M. (2016). Statistical learning of temporal community structure in the hippocampus. *Hippocampus*, 26(1), 3-8. doi: 10.1002/hipo.22523.

- Simon, H. A. (1955). On a class of skew distribution functions. *Biometrika*, 42(3-4), 425-440. Doi: 10.1093/biomet/4.
- Slobin, D. I. (1973). Cognitive prerequisites for the acquisition of grammar. In *Studies of child language development*, eds. C. A. Ferguson & D. I. Slobin, New York: Holt, Rinehart & Winston, 173-208.
- Sowa, J. F. (1987). Semantic networks. In *Encyclopedia of Artificial Intelligence*, ed. S. C. Shapiro, Hoboken: Wiley-Interscience, 1493-1511.
- Spielberger, C. D. & Starr, L. (1994). Curiosity and exploratory behavior. In *Motivation Theory and Research*. Hillsdale: Lawrence Erlbaum Associates.
- Spitzer, M. (1993). Associative networks, formal thought disorders and schizophrenia. On the experimental psychopathology of speech-dependent thought processes. *Nervenarzt*, 64(3), 147-59.
- Stanton, K. & Watson, D. (2017). Explicating the structure and relations of the Mood Disorder Questionnaire: Implications for screening for bipolar and related disorders. *Journal of Affective Disorders*, 220, 72-78. doi: 10.1016/j.jad.2017.05.046.
- Stella, M., Beckage, N. M., Brede, M. (2017). Multiplex lexical networks reveal patterns in early word acquisition in children. *Scientific Reports*, 7, 46730. doi: 10.1038/srep46730.
- Sternszus, R., Saroyan, A., Steinert, Y. (2017). Describing medical student curiosity across a four year curriculum: an exploratory study. *Medical Teacher*, 39(4), 377-382. doi: 10.1080/0142159X.2017.1290793.
- Steyvers, M. & Tenenbaum, J. B. (2005). The large-scale structure of semantic networks: statistical analyses and a model of semantic growth. *Cognitive Science*, 29(1), 41-78. doi: 10.1207/s15516709cog2901_3.
- Sutin, A. R., Beason-Held, L. L., Resnick, S. M., Costa, P. T. (2009). Sex differences in resting-state neural correlates of openness to experience among older adults. *Cerebral Cortex*, 19(12), 2797-802. doi: 10.1093/cercor/bhp066.
- Tripathi, R. K., Sarkate, P. V., Jalgaonkar, S. V., Rege, N. N. (2015). Development of active learning modules in pharmacology for small group teaching. *Education for Health (Abingdon)*, 28(1), 46-51. doi: 10.4103/1357-6283.161851.
- Tsimring, L. S. & Aranson, I. S. (1997). Localized and cellular patterns in a vibrated granular layer. *Physics Review Letters*, 79, 213. Doi: 10.1103/PhysRevLett.79.213.
- Vitevitch, M. S. (2008). What can graph theory tell us about word learning and lexical retrieval? *Journal of Speech Language and Hearing Research*, 51(2), 408-22. doi: 10.1044/1092-4388(2008/030).
- Vitevitch, M. S., Chan, K. Y., Goldstein, R. (2014). Insights into failed lexical retrieval from network science. *Cognitive Psychology*, 68 (Feb), 1-32. doi: 10.1016/j.cogpsych.2013.10.002.
- Volchenkov, D. & Blanchard, P. (2011). Fair and biased random walks on undirected graphs and related entropies. *Birkhäuser*, 380. doi: 10.1007/978-0-8176-4904-3_13.
- Voss, H. G. & Keller, H. (1983). *Curiosity and exploration: theories and results*. New York: Academic Press.
- Watts, D. J. & Strogatz, S. H. (1998). Collective dynamics of 'small-world' networks. *Nature*, 393(6684), 440-2. doi: 10.1038/30918.
- Webster, N. (1828). Webster's 1828 American Dictionary of the English Language. Foundation for American Christian Education; Facsimile of 1st edition, June 1, 1967.
- Widdows, D. (2004). *Geometry and meaning*. Stanford: CSLI Publications.
- Wymbs, N.F., Bassett, D.S., Mucha, P.J., Porter, M.A., Grafton, S.T. (2012). Differential recruitment of the sensorimotor putamen and frontoparietal cortex during motor chunking in humans. *Neuron*, 74(5), 936-46. doi: 10.1016/j.neuron.2012.03.038.

- Xue, G., Dong, Q., Chen, C., Lu, Z., Mumford, J. A., Poldrack, R. A. (2010). Greater neural pattern similarity across repetitions is associated with better memory. *Science*, 330(6000), 97-101. doi: 10.1126/science.1193125.
- Zhuang, K., Bi, M., Li, Y., Xia, Y., Guo, X., Chen, Q., Du, X., Wang, K., Wei, D., Yin, H., Qiu, J. (2017). A distinction between two instruments measuring dispositional mindfulness and the correlations between those measurements and the neuroanatomical structure. *Scientific Reports*, 7(1), 6252. doi: 10.1038/s41598-017-06599-w.
- Zurn, P. (forthcoming). Busybody, hunter, dancer: three historico-philosophical models of curiosity. In *Curious about curiosity: toward new philosophical explorations of the desire to know*, ed. Marianna Papastefanou, Cambridge: Cambridge Scholars Press.