

Edgework: Viewing curiosity as fundamentally relational

Perry Zurn⁵, Dale Zhou¹, David M. Lydon-Staley^{2,3,4}, and Danielle S. Bassett^{2,6,7,8,9,10,11}

¹Neuroscience Graduate Group, Perelman School of Medicine, University of Pennsylvania,
Philadelphia, PA 19104, USA

²Department of Bioengineering, School of Engineering and Applied Sciences, University of
Pennsylvania

³Annenberg School for Communication, University of Pennsylvania

⁴Leonard Davis Institute of Health Economics, University of Pennsylvania

⁵Department of Philosophy & Religion, American University, Washington, D.C.

⁶Department of Physics & Astronomy, College of Arts and Sciences, University of
Pennsylvania

⁷Department of Psychiatry, Perelman School of Medicine, University of Pennsylvania

⁸Department of Neurology, Perelman School of Medicine, University of Pennsylvania

⁹Department of Electrical & Systems Engineering, School of Engineering and Applied
Sciences, University of Pennsylvania

¹⁰Santa Fe Institute, Santa Fe, NM 87501 USA

¹¹To whom correspondence should be addressed: dsb@seas.upenn.edu

April 2, 2021

Abstract

Most theories of curiosity emphasize the acquisition of information. Such conceptualizations focus on the actions of the knower in seeking units of knowledge. Each unit is valued as an unknown and appropriated in becoming known. Yet, recent advances across a range of disciplines from philosophy to cognitive science suggest that it may be time to complement the acquisitional theory of curiosity with a connectional theory of curiosity. This alternative perspective focuses on the actions of the knower in seeking relations among informational units, laying down lines of intersection, and thereby building a scaffold or *network* of knowledge. Intuitively, curiosity becomes edgework. In this chapter, we dwell on the notion of edgework, wrestle with its relation to prior accounts, and exercise its unique features to craft alternative reasons for curiosity's value to humanity. To begin, we engage in a philosophical discussion of the evidence for connectional curiosity across the last two millennia in the Western intellectual tradition. We then move to a contemporary operationalization of connectional curiosity in the mathematical language of network science. To make our discussion more concrete, we walk through a case study of humans browsing Wikipedia. The groundwork laid, we turn to the practical question of how (if at all) the paradigm of curiosity as edgework manifests in the contemporary lives of humans today. Does such a conceptualization help us to better understand the relations between curiosity and mental health? Might the edgework paradigm explain the drive to build specific structures of knowledge? Would the account help us to encode, test, and validate existing theories of curiosity, or propose new ones? Could it clarify why and how our culture values curiosity, in its multiple manifestations, plethora of practices, and kindred kinds in many bodies? In considering interdisciplinary answers to these questions, we find that the notion of edgework offers a fresh, flexible, and explanatory account of curiosity. More broadly, it uncovers new opportunities to use the lens of science to examine, probe, and interrogate this important dimension of the human experience.

Introduction

The human mind is curious. It is strange, remarkable, and mystifying; it is eager, probing, questioning. Despite the pervasiveness of curiosity and its relevance for our well-being, many of us still pause in attempting to answer the fundamental question: What is curiosity? The quandary of how to define curiosity, its places, and its actions, is not a challenge that we newly face. Humanity has been attempting such definitions for centuries. In 1270, Aquinas suggests that curiosity is the desire to know [1]. Descartes in 1649 states that curiosity is a desire to understand [2]. In 1889, William James writes that curiosity is the impulse toward better cognition [3]. John Dewey in 1933 wrote that curiosity is an interest in problems provoked by the observation of things and the accumulation of material [4]. In 1994, Lowenstein suggests that curiosity is a feeling of deprivation produced by information gaps [5]. Even in modern science, Celeste Kidd writes that curiosity is a drive-state for information [6].

What is it precisely that we are desirous of, impulsive towards, interested in, deprived of, or driven towards? The answers – knowledge, information, solutions to problems, better cognition – are commonly thought of as items, which when acquired are satisfying. In fact, the analogy between hunger (thirst) and curiosity, as well as the analogy between food (drink) and information, underscores the acquisitional nature of our common conceptualizations [7]. Moreover, we think of curiosity as valuable to us because when we acquire the item of information, our life gets better: our uncertainty about the world is reduced [8], and we no longer feel deprived of information [5]. Upon reflection, the acquisitional account of curiosity seems intuitive. Yet, if we press the account, if we allow acquisitional actions of humans to play out upon the theatre of experience, we quickly come to realize that we are missing a key piece of the puzzle. An acquisition (or collection) of informational bits does not constitute knowledge.

Knowledge requires something more. Knowledge requires an understanding of the relations between bits of information: relations of cause, of correlation, and of consequence, to name a few. As Dewey writes, “[K]nowledge is a perception of those connections of an object which determine its applicability in a given situation. [...] An ideally perfect knowledge would represent such a network of interconnections that any past experience would offer a point of advantage from which to get at the problem presented in a new experience” [9]. Or consider the perspective of Henri Poincaré, when he writes “The aim of science is not things themselves, as the dogmatists in their simplicity imagine, but the relations among things; outside these relations there is no reality knowable” [10]. How might we expand the current acquisitional account of curiosity to an explicitly connectional account of curiosity? And if we did so, what sorts of affordances might such an account offer?

Here we frame curiosity as edgework. In doing so, we offer a justification, operationalization, investigation, and validation of a connectional account of curiosity. Our approach is explicitly interdisciplinary; we begin the discussion in philosophy, recast core notions in mathematics, and then exercise the ideas further in cognitive science, throughout drawing inspiration from physics and ethics as appropriate. The work we review ranges in inquisitory processes from laboratory experiments to mathematical models, and capitalizes on data of today and those of yesteryear. Upon this backdrop, we will see that the edgework account of curiosity provides new explanations for human behavior, and inspires a unique set of future investigative directions in understanding how we build knowledge of our networked world.

From Acquisitional to Connectional Curiosity

For centuries, even millennia, curiosity has been defined in the Western intellectual tradition as the desire to know [11]. Following suit, the contemporary fields of psychology and neuroscience typically understand curiosity to be the intrinsic motivation to seek information [6]. Implicit in these definitions is an acquisitional model of curiosity. Curiosity is the desire to know some thing or a behavior to seek some piece

of information. Curiosity is then quelled when knowledge is secured or information gaps are filled. But curiosity does not only acquire; it also, and perhaps more fundamentally, connects. It puts two and two together. It builds relations between ideas in the mind and between things in the world, steadily crafting and recrafting the network architecture of individual and collective knowledge. This is the connectional model of curiosity. And while this model has uniquely promising applications in network science today [12], it also has underappreciated roots in philosophies of yesteryear.

The acquisitional model of curiosity focuses, so to speak, on the nodes of knowledge to be obtained. This conceptualization has deep roots in the history of philosophy. Plutarch, for example, characterized the curious person as “passionate to find out” whatever is new or novel around town [13]. Augustine similarly describes the curious person as “lusting to see” everything from scientific marvels to religious wonders [14]. In the early modern period, this inquisitive passion became precisely the appetite for knowledge that promised to loose human society from its chains of ignorance [15]. Today, philosophers still conceptualize curiosity as the motivation to find out information or acquire specific epistemic goods [16] [17]. Importantly, the acquisitional model is rooted in units of knowers and units of knowledge: “I want to know X” or “I desire to find out X” [12]. On this model, the butterfly collector is paradigmatically curious because they aim to acquire elusive butterflies.

But where there are nodes, there are also edges. And, in fact, when the word curiosity (and its cognates in multiple languages) are tracked alongside words for “connection,” “link,” and “relation” within the Western intellectual tradition, another possibility appears: the connectional model. On this model, curiosity is edgework. As such, curiosity might collect information, track down answers, or imagine new possibilities, but it does so by building scaffolds or weaving webs. It builds connections, finds links, and follows threads. As Renaissance humanist Jean Luis Vives writes, curiosity tracks (and sometimes creates) the “threads” between things [18]. John Dewey similarly defines curiosity as the impulse to make the connections between things perceptible [9]. And that project will necessarily connect the curious person more deeply to their own interests and investments [19], but it will also, as William James astutely notes, necessarily “knit” new information together with the old [20]. On this model, it is not the collector but the butterfly conservationist who best embodies curiosity, for the latter aims to know the butterfly precisely within the relations that best sustain it.

While the traditional model of curiosity characterizes it as a desire to acquire new information, amassing knowledge in an ever-growing pile of nodes, the connectional model of curiosity characterizes it as a practice of building connections, whether between things to know or between knowers themselves. From a network science perspective, freeing curiosity from nodal gravity is priceless; it allows us to explore the curious energy of the edge.

Contemporary Operationalization

How might the paradigm of curiosity as edgework manifest in the contemporary lives of humans today? To address this question, it can be helpful to recast the core notions of edgework into a language that is currently used in modern science. The benefit of such a translation is that the assumptions, validity, and predictions of the paradigm can be tested in real data, using the central tool of science – inductive inference. In surveying a range of relevant options [21], it seems reasonable to begin by formalizing the notion of curiosity as connectional in the mathematical language of network science [22]. Here a network is a simplified representation of a complex system composed of units (nodes) and inter-unit relations (edges) that can be studied to understand system function [23]. As a simple concrete example, network science is commonly used to understand the spread of disease through a community, where the nodes are people and the edges are social contacts [24]. The pattern of edges among nodes explains how a virus can be transmitted through a population [25]. Whether applied to social, biological, technological, or physical systems, network science

captures the salient pattern of organization that is key to explaining how the system works.

In a connectional account of curiosity, the system we wish to study is knowledge and the system’s function is one of growth. Units of knowledge are represented as network nodes, and relations among units are represented as network edges [22]. Those edges can be weighted to indicate the strength of a relation among informational units, or to indicate a human’s likelihood to traverse that relation frequently in their thought patterns. Curious humans thus build networks by adding nodes and edges as they grow their body of knowledge. Using a more embodied description, curious minds engage in a kinesthetic practice, by walking through knowledge space and picking up relations to structure their thought architectures [26, 27]. The knowledge network inside each mind is therefore constantly evolving, and growing in manners and directions that are idiosyncratic. We walk along the links that connect bits of information, and then stand at the void, and shout across. With Walt Whitman, we “think a thought of the clef of the universe, and of the future” [28]. When we find or discover a unit of information that fully or partially fills the void, we have built a new bit of scaffold in our knowledge network.

Case Study of Curiosity as Edgework

How does this walking, shouting, and filling happen in everyday life? We recently set out to watch humans walk on collective knowledge networks (Wikipedia), thereby building their own individual knowledge networks [29]. After providing informed consent for the study, participants browsed Wikipedia for 15 minutes a day for 21 days. While they browsed, a piece of software installed on their computer provided us with information about which sites they visited, in what order, and for how long. The data were then analyzed by treating each Wikipedia page as a node in a network, and by treating the transition from one page to another as a step along the edge connecting them [30, 31]. We defined long steps to be those taken between pages with dissimilar content, and short steps to be those taken between pages with similar content. The sequence of steps that each person took reflected their characteristic style of curiosity, and thus of knowledge network building.

In the many hours of data collected, we saw clear evidence for two archetypal styles of curiosity: the busybody and the hunter [27]. The hunter “wishes [they] had a few hundred helpers and good, well-trained hounds that [they] could drive into the history of the human soul to round up [their] game” (Nietzsche, 1886) [32] in a targeted search for information; in contrast, the busybody will “frisk about, and rove about, at random, wherever they please” (Philo Judaeus, 25 BC) [33]. These kinesthetic signatures of the hunter and busybody naturally lead to the creation of different individualized knowledge networks. Consistent with this expectation, we found that the hunters in the participant group took walks that traced out networks with high clustering, meaning that they navigated among sets of 3 or more pages with high similarity, taking short steps in a local neighborhood [34]. In contrast, the busybodies in the participant group took walks that traced out networks with low clustering and long path length, meaning that they navigated among pages that were dissimilar, taking long steps spanning many distinct neighborhoods.

Interestingly, one personality trait stood out in explaining the different network building practices of the busybodies and hunters: their sensitivity to information deprivation [35, 5, 36]. In a separate testing session, participants had completed a scale measuring their deprivation sensitivity by determining the degree to which the following statements characterized them: “Thinking about solutions to difficult conceptual problems can keep me awake at night”; “I can spend hours on a single problem because I just can’t rest without knowing the answer”; “I feel frustrated if I can’t figure out the solution to a problem, so I work even harder to solve it”; “I work relentlessly at problems that I feel must be solved”; “It frustrates me not having all the information I need” [37]. Individuals high in deprivation sensitivity have a drive to eliminate the unknown as they encounter new information and recognize gaps in their knowledge. Consistent with this trait, we found that individuals high in deprivation sensitivity were hunters building densely clustered

networks, while those low in deprivation sensitivity were busybodies building sparse and loopy networks. In other words, individual differences in deprivation sensitivity lead to the creation of knowledge networks with distinct architectures.

Curiosity as edgework in a changing mental world

Does our mental state explain daily variation in the connectional knowledge structures that we build?

Beyond trait deprivation sensitivity, is there any evidence for mental state to affect the knowledge network structures we build? There is a rich literature that considers the multiple timescales over which the experience of curiosity unfolds [38]. In addition to investigations of curiosity that focus on between-person differences in the tendency to consistently experience high levels of curiosity across periods of time on the order of months and years [37, 39], curiosity can also be a transient experience, lasting only moments [40, 5]. We can additionally consider experiences of curiosity at intermediate timescales on the order of days [41, 39, 42]. These shorter-term fluctuations in curiosity are intimately tied to dynamic mental states, in particular to the waxing and waning of our emotional experiences [43, 44, 45]. Positive affect facilitates an openness to experiences, even challenging ones, that is at the core of many definitions of curiosity (e.g., “I view challenging situations as opportunities to grow and learn”, “I frequently seek out opportunities to challenge myself and grow as a person” [46]). By broadening individuals’ behaviors and cognitive repertoires, positive affect states facilitate engagement in diverse and novel courses of action [47]. Negative emotions, in contrast, can promote a narrowing of behavior directed towards the immediate situation [48].

Extending this consideration of multiple timescales and dynamics in self-reported states of curiosity and mental states to edgework, we decomposed the 21 days of knowledge network building on Wikipedia into three, one-week periods. Performing this decomposition revealed substantial week-by-week fluctuations in the edgework being performed by participants and also indicated that changing mental states explained these fluctuations. We found, for example, that weeks of higher than usual sensation-seeking [49], weeks when participants had relatively greater craving for excitement and new experiences than typical for them [50], were also weeks during which participants browsed with longer steps than usual, i.e., traversing more dissimilar Wikipedia pages and building more spindly, diffuse networks [29]. Further, our unpublished data show that subclinical levels of depression are associated with the tendency to walk more locally over self-similar inter-page links; participants who were more depressed favored building close-by, whereas participants who were less depressed favored laying down far-flung edges. This narrowing of information search in the context of depressed affective states parallels the narrowing of self-reported curiosity when experiencing negative affect [48]. Yet, we note that, in the connectional curiosity case, we view this decrease in the tendency to seek diverse and novel information in the context of negative affect as a modulation of the type of edgework that is practiced rather than a reduction in curiosity.

Our observations that styles of edgework fluctuate across time alongside naturalistic fluctuations in mental state highlight the important role that our mental states play in the types of knowledge structures we build. Understanding the time-varying nature of edgework more deeply will require the acquisition of fine-grained measurements of each person’s time-varying mental state concurrently with intensive assessment of edgework. Such investigations will benefit from increasingly feasible approaches to capturing dynamics in both mental state [51, 52] and media engagement [53].

Curiosity as edgework on a complex external landscape

How does the way we walk in solitude relate to the way we engage with information in society?

So when we seek information, we each seek it differently, thereby constructing fundamentally different networks of knowledge. Do the principles of information seeking relate to the principles of information receiving, and if so how? It seems plausible that a person would learn best when they are presented with information that has a network structure that is consistent with that which they would intrinsically prefer to build in solitude. In other words, if we present information in a loopy manner to a person who builds dense rigid networks, they might be slower to learn. Or, if we present information in a dense rigid manner to a person who builds sparse loopy networks, they might be slower to learn. Alternatively, might some discrepancy between the presented architecture and the sought architecture be advantageous for learning? It seems plausible that the novelty of an unsought-for architecture could raise the learner’s level of arousal and increase their attention. Are either (or both) of these competing intuitions correct?

To answer this question, one would need an experimental paradigm to determine how well humans learn knowledge presented with a given network architecture. Recently, such paradigms have begun to be developed in the field of *graph learning*, which stems from early advances in statistical learning and the study of artificial grammars [54]. In both earlier areas of inquiry, seminal studies investigated the human capacity to learn the transition probabilities between two sequential events [55]; because these events occurred next to one another in time the respective transition probabilities were referred to as *adjacent* probabilities [56]. More recently, the field of graph learning presses beyond adjacent to non-adjacent transition probabilities, seeking to understand how humans learn and represent the local, mesoscale, and macroscale architecture of networks in the world around them [57, 54]. A key focus lies in how humans build network models as they simply experience sequences of words, actions, images, or events [58, 59, 60].

In a typical graph learning experiment, the participants will see a sequence of informational bits. The transition probabilities between informational bits obeys a pre-defined network structure that is never directly revealed to the participant. The question is how well the participant can infer – or accurately learn – this network over time. Learning is typically measured by how quickly humans react to each informational bit; did they recognize and/or anticipate that bit, or detect ways in which that bit might have been minutely changed? These sorts of reactions, and how quickly they happen, can be used to index how much that part of the network has been learned. Data suggest that most humans learn new information best when it is presented in a modular network organization [60, 61]: where some informational units are densely interconnected (and thus often occur in close temporal proximity), forming clusters or modules. Yet, there is also marked variability in how easily humans learn distinct topologies [60, 59]. These individual differences are precisely the features that are needed to determine whether (and how) a human’s preferences to engage with a network topology in a learning experiment are aligned with (or divergent from) that same human’s preferences to build a network topology in a solitary curiosity-driven seeking process. Future studies could recruit participants to engage in both types of tasks to assess how structural preferences constrain both intrinsically motivated search and receptivity to extrinsically offered information.

Curiosity as edgework reshapes previous theoretical constructs

Does connectional curiosity validate the information gap or compression progress theories of curiosity?

Framing curiosity as edgework does not only raise new opportunities to study the internal landscape of mental experience and the external landscape of societally-embedded learning. It also fundamentally reshapes the theoretical discourse, impacting the very roots of our inquiry. Many have acknowledged that it is not easy to precisely define curiosity or to build an appropriate theoretical construct to parameterize

its compositional concepts. Throughout the history of modern science, scholars have engaged a range of approaches that seek to explain curiosity’s manifestations in human and non-human animals. Prior theories often do not define curiosity as a construct in and of itself, but as being composed of inextricable interactions among other constructs, such as those of learning, attention, reward, uncertainty, and exploration [6, 62]. According to these theories, curiosity is an emotion or drive defined by reward or uncertainty [5, 63, 64, 6] and is explained by its ability to compel learning [65, 66, 64], attention [67, 68], and exploration [69, 70]. Yet, these explanations beg the question of whether (and how) curiosity is distinct from other constructs, being its own process rather than motivating other processes.

One particularly appealing answer to this question is provided by the notion of curiosity as edgework. By framing curiosity as a building process, the compositional concepts become clear: a system of knowledge composed of relations among informational units, which grows over time by the process of relation-finding actions [22, 71]. By operationalizing those compositional concepts in the mathematical language of network science, the field has the opportunity to use computational science to test two key prior theories of curiosity: the information gap theory and the compression progress theory [5, 72]. In the former, curiosity is the drive to fill an information gap, to find the piece of knowledge that is currently missing in a mental model [5, 63]. In the latter, curiosity is the drive to seek information that improves the compression of a person’s mental model of the world or a sector thereof, enhancing the compactness, efficiency, and flexible use of knowledge [72, 71].

These two theories have proven difficult to validate definitively. Progress has in part been stymied by difficulties in describing the two key concepts (an information gap and a progress of compression) in a way that allows us to measure them explicitly in contemporary experimental paradigms. The translation from theoretical concept to measurable variable is a perennial challenge, which can sometimes be addressed by choosing an appropriate mathematical model, whose variables are at once measurable and representative of the relevant concepts. What empirical and computational data could validate either or both frameworks? By framing curiosity as edgework and operationalizing edgework as a network building process, we gain traction in translating theoretical concepts to measurable variables. First, an information gap can be explicitly defined as a cavity in the network, which is the absence of edges in a portion of the network forming a hole between and among concepts. Notably, a network cavity can now be detected algorithmically using recent advances in applied algebraic topology [73, 74], a subfield of mathematics. Similarly, when curiosity is a network building process, then compression progress can be quantified as a change in the network’s compressibility, which can now be calculated using recent work at the intersection of rate-distortion theory and information theory [75].

Each theory predicates distinct observations of the building process. The growth of individual and collective knowledge networks accompanied by a closing of cavities would support the validity of information gap theory; initial evidence for such closing of cavities arises in the development of language in toddlers [73] and collective knowledge in Wikipedia [74]. Knowledge network growth accompanied by an increase in compressibility would support the validity of compression progress theory; evidence for such an increase has yet to accrue, pending direct assessments of network compressibility. The framing of curiosity as edgework does not only provide a way of validating a given theory, but also a way of testing whether the two theories are compatible or mutually exclusive. Tighter networks are created by filling in gaps, and some sorts of such gap-less networks are markedly compressible [75]. Looser networks are created by leaving gaps open, and some sorts of such gap-full networks can be less compressible. It would be of interest in future work to determine the relative rates of gap-filling in the growth of individual versus collective knowledge networks, and to determine whether those rates are aligned with (or divergent from) changes in the networks’ compressibility. Such an investigation could elucidate whether the two theories are describing the same generative process, or distinct generative processes that can be engaged either synergistically or competitively.

Curiosity as edgework poses a novel theoretical construct

Does connectional curiosity inherently pose an alternative description of curiosity’s utility to the individual human?

In addition to clarifying and reshaping prior theories of curiosity, the framing of edgework predicates its own distinct explanations. To see how these predictions arise, it is important to think more deeply about the nature of knowledge. As we have seen, knowledge can be usefully modeled as a network of concepts and their relations [76]. Some of those relations are fixed and immutable such as, for example, the relations among whole numbers. Others are flexible in the sense that they can be altered by context or by the addition, deletion, or transmutation of concepts around them [77]. How does the structure of a network determine its flexibility? Consider an architect constructing the scaffolding of a new building, an engineer constructing the skeleton of a robotic system [78], or natural selection evolving the mechanics of a living animal [79]. The stability, longevity, and rigidity of such structures is made possible by the geometry of their linkages, as Scottish physicist James Clerk Maxwell pointed out in the middle of the 1800s [80]. A triangle made from steel bars connected by hinges is rigid because pressure on any bar does not change the triangle’s shape. Conversely, a square made from the same materials is flexible because pressure applied to any bar will collapse the square into a rhombus. By combining rigid sections with non-rigid sections, larger structures can be built that marry flexibility and constraints in a single kinematics.

As we build our knowledge networks through curious practice, do we tether concepts lightly? Or do we perhaps leave some relatively untethered, to allow for their movement in the winds of our future experience? Do we combine sections of rigid (triangular) conceptual relations with sections of flexible (squarish) conceptual relations, allowing our minds to bend as and when necessary? It seems reasonable to consider that curiosity’s utility to a single human might lie in its ability to guide our knowledge network formation to maintain flexibility and eschew rigidity. The relation-seeking actions of curiosity may continuously result in a structure that, like macromolecules, is amenable to conformational change [81]. In isolation, the process of filling information gaps (closing triangles in the network) will create a rigid knowledge, whereas the process of reaching out to the unknown (adding otherwise untethered concepts) will create a knowledge that can undergo shape-morphing. Is it perhaps by this flexibility that we gain insight, the “Aha!” moment as two seemingly distant concepts come close enough in our conceptual space to touch?

Such questions can be operationalized in a formal mathematical model, with quantifiable measures of the conformability of a growing network. Intuitively, one can take the design principles for conformational change recently developed for physical spaces [82, 83], and transform them into conceptual spaces. Building on the study of linkages and transformable metamaterials [84, 78, 85], one can calculate the degrees of freedom of growing knowledge networks, reflecting the flexibility of a set of related concepts in d dimensional concept space. One can also determine the states of self stress of growing knowledge networks, or the subsets of related concepts that tend to maintain their inter-relations under some perturbation [80, 86]. Such calculable, theory-based measures can be identically applied to (i) experimental data acquired from information seeking experiments in single humans, and (ii) archived data from collective information seeking in groups of humans. Moreover, they can be used to assess how knowledge networks grow over development, as children mature into adulthood. The fact that looser and more fluid, as opposed to tighter and more crystalline, knowledge networks may be easier to reconfigure, could explain why children are able to better learn and flexibly use information from abstract schema than adults [70]. Evidence from neuroscience further suggests that the patterns of inter-regional connections in the human brain may increasingly prioritize less compressed transmission (less abstract) as development occurs [87], potentially altering the sorts of inter-conceptual connections made and carried with a mind into later life.

Curiosity as edgework embodies the edgework of the brain

How does the mind curious-ing build a network of knowledge using a network of neurons?

Considering curiosity as edgework situates the information seeker as modifying and interacting with its internal and external information environment to build a model of the world [71]. Behaviors that selectively deploy attention rely on the capacity of the brain to process information from the external environment and to transform that information within the internal environment of neural representations [67]. With the undertaking of edgework, information seekers build network models of information that can be abstractly relational, social, and rewarding [88, 89, 90, 60, 91, 59, 92]. Recent efforts have investigated the capacity to process information from the environment and within the brain using novel approaches in graph learning [93, 76, 22, 57, 94]. Ongoing efforts seek to understand where in the brain we represent each node and edge in the network, as well as where (and how) we encode the network’s mesoscale and global architecture. How are these representations and encodings dynamically created, modified, and (sometimes) forgotten [95, 87]? Following Dewey, we might ask how “such a network of interconnections” is employed to “offer a point of advantage from which to get at the problem presented in a new experience” [9]. Answering these questions requires us to understand not only the locations in the brain where the world leaves an impression (like a foot in wet sand), but also how those impressions relate to one another reflecting the trajectory of information (like the series of footprints in the sand reflecting the trajectory of the walker from beach to cove).

Understanding how brain regions and their respective encodings relate to one another, and share information with one another, is the purview of the nascent field of network neuroscience [93]. Here, brain regions form network nodes, and inter-regional relations or interactions form network edges. Network neuroscience defines mathematical models and metrics that characterize the neural efficiency of information transfer and the spatial distribution of information gathered both from within the brain and from the external environment [93, 96]. The approach reveals that the human brain network is characterized by a striking gradient of computations organized along a principal axis spanning the cortex [97], from unimodal sensation to transmodal higher-order abstract processing. At the transmodal extreme of the gradient, localized regions in the frontoparietal, default-mode, and limbic systems have been associated with learning, abstraction, self-referential processing, semantic cognition, emotional regulation, uncertainty monitoring, and reward [97, 98]. In so far as curiosity has been traditionally defined as inextricable from some of these cognitive functions, it is therefore unsurprising that prior work has also found associations between curiosity-driven behavior and the neural activity of specific regions within these subnetworks [63, 64, 99]. However, it remains unclear how localized processes within separate brain regions or systems are fused together within the wider brain network, as prior theories necessitate.

Curiosity as edgework re-frames this question by asking: how does the network of the human brain communicate the consequence of local computations (supporting processes ranging from sensation, emotion, and cognition) in order to build a model situating oneself in relation to the world? One possible answer lies in a structural isomorphism between the brain and the world, which supports the development of internal mental models. It is now clear that the network organization that efficiently transmits and flexibly represents information is similar for both the information processed by the human brain [87] and the information processed during human behavior [57, 61, 75]. Both the brain and behavior rely on the usage of hubs and hierarchical organization. Hubs are central landmarks that are highly connected to other units. Hierarchical organization is a hallmark of efficient spatial nesting of modules within modules and the formation of abstract categories of knowledge [71, 100]. This interesting isomorphism between the edgework of brains and behavior could arise from shared evolutionary and developmental pressures for efficient yet flexible spatiotemporal information processing and representation [71, 54]. Shared constraints of efficiency and flexibility on the brain and behavior may shape the capacity to learn statistical relations, form a system of semantic categories, deploy attention, and explore new information [100, 58, 67, 101, 70]. When we seek information, the ability of an external source to transmit information is intertwined with the ability to transmit and dynamically transform internal representations of information within the brain network [61, 87, 74]. In doing so, the brain network

can shape and modify information to send only behaviorally relevant inputs to the hippocampal-entorhinal circuit for representation of statistical regularities, prediction, and planning [88, 102, 89, 90].

Curiosity as edgework is valued

How is connective curiosity valued in our society today?

How does our culture value curiosity? Does it only and ever value the acquisition of knowledge? Or is there evidence that our culture values the laying down of connections, and of some connections more than others? Two recent studies provide evidence for the latter, both in the recognition of scientific discoveries, and in the evaluation of classroom material. The first study reports that science progresses as much through identifying uncharted gaps as through advancing solutions within scientific communities; the latter, filling-the-gap (or tessellating-the-cavity) discoveries are more frequently awarded Nobel prizes than discoveries adding edges at the boundaries of collective knowledge [74]. The second study reports that college-level mathematics textbooks are more likely to be rated highly on GoodReads when they progressively fill (rather than leave open) gaps in the knowledge network by connecting concepts introduced early with concepts introduced late [103]. Together, these studies suggest that there exist some domains of human inquiry where the act of placing a connection that fills a network cavity is more greatly valued than the act of placing a connection to the unknown at the frontiers of knowledge.

Is this same valuation strategy used in other aspects of our society? Consider, for example, the choice of whose ideas to discuss, as well as how (and whether) to attribute those ideas to the person who created them [104]. Both the choice and the attribution are processes of valuation, and occur in verbal communication, social and mass media, and written documents [105]. In scientific articles and classroom textbooks, for example, new material references older material, offering a record of the questions that have been and are being asked, and the answers of greatest import and value [106]. Are articles or chapters that fill a network cavity in the space of scientific knowledge cited more frequently than those that add a new connection at the frontiers of that space? The field is poised to answer this question. And what makes this area of inquiry particularly exciting is the fact that it naturally interdigitates with the field of ethics: namely, the ethics of knowing, addressing epistemic justice and injustice [107]. Citation ethics seeks to understand and address existing biases in citation practices, including the bias against women and racial/ethnic minorities [108, 109, 110, 111, 112, 113]. Determining how we value the edges that knowers newly place will require both an understanding of the network placements we value (those tessellating or those expanding) and the placers we value (those privileged or those marginalized). Such an investigation could shed important new light on the factors that determine how our culture values curiosity (and curious minds embodied in gendered and racialized beings), simultaneously informing the development of an explicit and systematic ethics of citation practices germane to classroom learning and the scientific enterprise [114, 115].

Conclusion

In being curious, we connect our minds to previously known knowns and to knowns just discovered. In acting curiously, we connect our bodies to previously existing knowers and to knowers alive today. We follow threads of ideas, weave lines of thought, and weld hinges among intersecting concepts. We create a scaffold of knowledge with the wrench of attention, the hammer of reward, the spirit level of exploration, the pliers of learning, and the screwdriver of uncertainty. Some of us shy away from the world, choosing instead to tinker in the bowels of the scaffold; others of us scale to the top and peer out at the world from above, or scramble to the edges and peek at the world from the side. We see a world of possibilities ... a world of what could be when we exchange acquisition for connection, capture for construction, and conquest for relation. Curiosity as edgework offers a new lens into the constructs, affordances, and values of seeking communion

with our world.

Acknowledgements

We are thankful for the insightful feedback and comments from Jason Z. Kim, Shubhankar Patankar, and Pixel Xia. We gratefully acknowledge support from the John D. and Catherine T. MacArthur Foundation, the Alfred P. Sloan Foundation, the ISI Foundation, and the Center for Curiosity. The content is solely the responsibility of the authors and does not necessarily represent the official views of any of the funding agencies.

Citation Diversity Statement

Recent work in neuroscience and other fields has identified a bias in citation practices such that papers from women and other minorities are under-cited relative to the number of such papers in the field [113, 111, 108, 110, 109]. Here we sought to proactively consider choosing references that reflect the diversity of the field in thought, form of contribution, gender, and other factors. First, we obtained predicted gender of the first and last author of each reference by using databases that store the probability of a name being carried by a woman [109, 116]. By this measure (and excluding self-citations to the first and last authors of our current paper), our references contain 19.28% woman/woman, 16.1% woman/man, 8.14% man/woman, and 56.48% man/man. This method is limited in that a) names, pronouns, and social media profiles used to construct the databases may not, in every case, be indicative of gender identity and b) it cannot account for intersex, non-binary, or transgender people. We look forward to future work that could help us to better understand how to support equitable practices in science. Second, we obtained predicted racial/ethnic category of the first and last author of each reference by databases that store the probability of a first and last name being carried by an author of color [117, 118]. By this measure (and excluding self-citations), our references contain 7.32% author of color (first)/author of color(last), 16.21% white author/author of color, 14.24% author of color/white author, and 62.23% white author/white author. This method is limited in that a) names and Florida Voter Data to make the predictions may not be indicative of racial/ethnic identity, and b) it cannot account for Indigenous and mixed-race authors, or those who may face differential biases due to the ambiguous racialization or ethnicization of their names. We look forward to future work that could help us to better understand how to support equitable practices in science.

References

- [1] T. Aquinas, *Summa Theologica*. Notre Dame: Christian Classics, 1981.
- [2] R. Descartes, “Passions of the soul,” in *The Philosophical Writings of Rene Descartes* (J. Cottingham, ed.), p. sec. 88, Cambridge: Cambridge University Press, 1985, 1649.
- [3] W. James, *Talks to teachers on psychology: And to students on some of life’s ideals*. Henry Holt and Company, 1925.
- [4] J. Dewey, *How we think*. D. C. Heath and Company, 1933.
- [5] G. Loewenstein, “The psychology of curiosity: A review and reinterpretation.,” *Psychological bulletin*, vol. 116, no. 1, p. 75, 1994.
- [6] C. Kidd and B. Y. Hayden, “The psychology and neuroscience of curiosity,” *Neuron*, vol. 88, no. 3, pp. 449–460, 2015.
- [7] *Cofer, C N and Appley, M H*. Wiley, New York, 1964.
- [8] J. D. Cohen, S. M. McClure, and A. J. Yu, “Should i stay or should I go? How the human brain manages the trade-off between exploitation and exploration,” *Philos Trans R Soc Lond B Biol Sci*, vol. 362, no. 1481, pp. 933–942, 2007.
- [9] J. Dewey, *Democracy and Education*. Simon and Brown, New York, 1916.
- [10] H. Poincare, “Science and hypothesis,” pp. xxi – xxvii, 1905.
- [11] P. Zurn, *Curiosity and Power: The Politics of Inquiry*. University of Minnesota Press, 2021.
- [12] D. S. Bassett and P. Zurn, *Curious Minds*. MIT Press, forthcoming 2022.
- [13] Plutarch, “On being a busybody,” 100;2005.
- [14] S. Augustine, “Confessions, trans,” *W. Watts (London, 1631)*, 1961.
- [15] B. Benedict, *Curiosity: A Cultural History of Early Modern Inquiry*. University of Chicago Press, 2003.
- [16] I. Inan, *The Philosophy of Curiosity*. Routledge, 2011.
- [17] L. Watson, “Curiosity and inquisitiveness,” 2018.
- [18] J. L. Vives, *On Education*. Cambridge University Press, 1531; 1913.
- [19] J.-J. Rousseau, *Emile*. Appleton and Company, 1762; 1909.
- [20] W. James, *The Principles of Psychology, vl. 2*. Dover Publications, 1890; 1918.
- [21] L. Torres, A. S. Blevins, D. S. Bassett, and T. Eliassi-Rad, “The why, how, and when of representations for complex systems,” *arXiv*, vol. 2006, p. 02870.
- [22] D. S. Bassett, “A network science of the practice of curiosity,” in *Curiosity Studies: A new ecology of knowledge* (P. Zurn and A. Shankar, eds.), Minnesota Press, 2020.
- [23] M. Mitchell, *Complexity: A Guided Tour*. Oxford University Press, 2011.
- [24] S. Bansal, J. Read, B. Pourbohloul, and L. A. Meyers, “The dynamic nature of contact networks in infectious disease epidemiology,” *J Biol Dyn*, vol. 4, no. 5, pp. 478–489, 2010.

- [25] G. Pullano, L. D. Domenico, C. E. Sabbatini, E. Valdano, C. Turbelin, M. Debin, C. Guerrisi, C. Kengne-Kuetché, C. Souty, T. Hanslik, T. Blanchon, P.-Y. Boëlle, J. Figoni, S. Vaux, C. Campèse, S. Bernard-Stoecklin, and V. Colizza, “Underdetection of COVID-19 cases in France threatens epidemic control,” *Nature*, vol. Epub Ahead of Print, 2020.
- [26] P. Zurn and D. S. Bassett, “On curiosity: a fundamental aspect of personality, a practice of network growth,” *Personality Neuroscience*, vol. 1, 2018.
- [27] P. Zurn, “Chapter 2 busybody, hunter, and dancer: Three historical models of curiosity,” *Toward New Philosophical Explorations of the Epistemic Desire to Know: Just Curious About Curiosity*, p. 26, 2019.
- [28] W. Whitman, “On the beach at night alone,” in *Leaves of Grass*, James R. Osgood and Company (Boston), 1881.
- [29] D. M. Lydon-Staley, D. Zhou, A. S. Blevins, P. Zurn, and D. S. Bassett, “Hunters, busybodies, and the knowledge network building associated with curiosity,” 2019.
- [30] I. Scholtes, “When is a network a network? multi-order graphical model selection in pathways and temporal networks,” in *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*, pp. 1037–1046, 2017.
- [31] R. West and J. Leskovec, “Human wayfinding in information networks,” in *Proceedings of the 21st international conference on World Wide Web*, pp. 619–628, 2012.
- [32] F. W. Nietzsche, *Philosophy in the Tragic Age of the Greeks*, vol. 6075. Regnery publishing, 1996.
- [33] C. D. Yonge *et al.*, *The works of Philo Judaeus, the contemporary of Josephus*, vol. 1. Bohn, 1854.
- [34] J.-P. Onnela, J. Saramäki, J. Kertész, and K. Kaski, “Intensity and coherence of motifs in weighted complex networks,” *Physical Review E*, vol. 71, no. 6, p. 065103, 2005.
- [35] J. A. Litman, “Interest and deprivation factors of epistemic curiosity,” *Personality and individual differences*, vol. 44, no. 7, pp. 1585–1595, 2008.
- [36] J. A. Litman and T. L. Jimerson, “The measurement of curiosity as a feeling of deprivation,” *Journal of personality assessment*, vol. 82, no. 2, pp. 147–157, 2004.
- [37] T. B. Kashdan, M. C. Stikma, D. J. Disabato, P. E. McKnight, J. Bekier, J. Kaji, and R. Lazarus, “The five-dimensional curiosity scale: Capturing the bandwidth of curiosity and identifying four unique subgroups of curious people,” *Journal of Research in Personality*, vol. 73, pp. 130–149, 2018.
- [38] R. Pekrun, “The murky distinction between curiosity and interest: State of the art and future prospects,” *Educational Psychology Review*, vol. 31, no. 4, pp. 905–914, 2019.
- [39] D. M. Lydon-Staley, P. Zurn, and D. S. Bassett, “Within-person variability in curiosity during daily life and associations with well-being,” *Journal of personality*, 2019.
- [40] D. E. Berlyne, “Conflict, arousal, and curiosity,” 1960.
- [41] E. Garrosa, L. M. Blanco-Donoso, I. Carmona-Cobo, and B. Moreno-Jiménez, “How do curiosity, meaning in life, and search for meaning predict college students’ daily emotional exhaustion and engagement?,” *Journal of Happiness Studies*, vol. 18, no. 1, pp. 17–40, 2017.
- [42] T. B. Kashdan and M. F. Steger, “Curiosity and pathways to well-being and meaning in life: Traits, states, and everyday behaviors,” *Motivation and Emotion*, vol. 31, no. 3, pp. 159–173, 2007.
- [43] E. Diener and C. Diener, “Most people are happy,” *Psychological science*, vol. 7, no. 3, pp. 181–185, 1996.

- [44] J. R. Rodrigue, K. R. Olson, and R. P. Markley, “Induced mood and curiosity,” *Cognitive Therapy and Research*, vol. 11, no. 1, pp. 101–106, 1987.
- [45] T. G. Reio and J. L. Callahan, “Affect, curiosity, and socialization-related learning: A path analysis of antecedents to job performance,” *Journal of Business and Psychology*, vol. 19, no. 1, pp. 3–22, 2004.
- [46] T. B. Kashdan, M. W. Gallagher, P. J. Silvia, B. P. Winterstein, W. E. Breen, D. Terhar, and M. F. Steger, “The curiosity and exploration inventory-ii: Development, factor structure, and psychometrics,” *Journal of research in personality*, vol. 43, no. 6, pp. 987–998, 2009.
- [47] B. L. Fredrickson, “The broaden-and-build theory of positive emotions,” *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, vol. 359, no. 1449, pp. 1367–1377, 2004.
- [48] B. L. Fredrickson and C. Branigan, “Positive emotions broaden the scope of attention and thought-action repertoires,” *Cognition & emotion*, vol. 19, no. 3, pp. 313–332, 2005.
- [49] M. Zuckerman, *Behavioral expressions and biosocial bases of sensation seeking*. Cambridge university press, 1994.
- [50] D. M. Lydon-Staley, E. B. Falk, and D. S. Bassett, “Within-person variability in sensation-seeking during daily life: Positive associations with alcohol use and self-defined risky behaviors,” *Psychology of addictive behaviors*, 2019.
- [51] R. Larson and M. Csikszentmihalyi, “The experience sampling method,” in *Flow and the foundations of positive psychology*, pp. 21–34, Springer, 2014.
- [52] D. M. Lydon-Staley, I. Barnett, T. D. Satterthwaite, and D. S. Bassett, “Digital phenotyping for psychiatry: Accommodating data and theory with network science methodologies,” *Current opinion in biomedical engineering*, vol. 9, pp. 8–13, 2019.
- [53] B. Reeves, N. Ram, T. N. Robinson, J. J. Cummings, C. L. Giles, J. Pan, A. Chiatti, M. Cho, K. Roehrick, X. Yang, *et al.*, “Screenomics: A framework to capture and analyze personal life experiences and the ways that technology shapes them,” *Human-Computer Interaction*, pp. 1–52, 2019.
- [54] C. W. Lynn and D. S. Bassett, “How humans learn and represent networks,” *Proc Natl Acad Sci U S A*, vol. 117, no. 47, pp. 29407–29415, 2020.
- [55] J. R. Saffran, R. N. Aslin, and E. L. Newport, “Statistical learning by 8-month-old infants,” *Science*, vol. 274, no. 5294, pp. 1926–1928, 1996.
- [56] A. R. Romberg and J. R. Saffran, “Statistical learning and language acquisition,” *Wiley Interdiscip Rev Cogn Sci*, vol. 1, no. 6, pp. 906–914, 2010.
- [57] E. A. Karuza, S. L. Thompson-Schill, and D. S. Bassett, “Local patterns to global architectures: influences of network topology on human learning,” *Trends in cognitive sciences*, vol. 20, no. 8, pp. 629–640, 2016.
- [58] A. C. Schapiro, T. T. Rogers, N. I. Cordova, N. B. Turk-Browne, and M. M. Botvinick, “Neural representations of events arise from temporal community structure,” *Nature neuroscience*, vol. 16, no. 4, pp. 486–492, 2013.
- [59] A. E. Kahn, E. A. Karuza, J. M. Vettel, and D. S. Bassett, “Network constraints on learnability of probabilistic motor sequences,” *Nature human behaviour*, vol. 2, no. 12, pp. 936–947, 2018.
- [60] C. W. Lynn, A. E. Kahn, and D. S. Bassett, “Abstract representations of events arise from mental errors in learning and memory,” *Nature Communications*, vol. 11, no. 1, p. 2313, 2020.

- [61] C. W. Lynn, L. Papadopoulos, A. E. Kahn, and D. S. Bassett, “Human information processing in complex networks,” *Nature Physics*, vol. In Press, 2020.
- [62] J. Gottlieb and P.-Y. Oudeyer, “Towards a neuroscience of active sampling and curiosity,” *Nature Reviews Neuroscience*, vol. 19, no. 12, pp. 758–770, 2018.
- [63] M. J. Kang, M. Hsu, I. M. Krajchich, G. Loewenstein, S. M. McClure, J. T.-y. Wang, and C. F. Camerer, “The wick in the candle of learning: Epistemic curiosity activates reward circuitry and enhances memory,” *Psychological science*, vol. 20, no. 8, pp. 963–973, 2009.
- [64] M. J. Gruber, B. D. Gelman, and C. Ranganath, “States of curiosity modulate hippocampus-dependent learning via the dopaminergic circuit,” *Neuron*, vol. 84, no. 2, pp. 486–496, 2014.
- [65] P.-Y. Oudeyer, J. Gottlieb, and M. Lopes, “Intrinsic motivation, curiosity, and learning: Theory and applications in educational technologies,” in *Progress in brain research*, vol. 229, pp. 257–284, Elsevier, 2016.
- [66] P.-Y. Oudeyer, “Computational theories of curiosity-driven learning,” 2018.
- [67] J. Gottlieb, “Attention, learning, and the value of information,” *Neuron*, vol. 76, no. 2, pp. 281–295, 2012.
- [68] J. Gottlieb, M. Hayhoe, O. Hikosaka, and A. Rangel, “Attention, reward, and information seeking,” *Journal of Neuroscience*, vol. 34, no. 46, pp. 15497–15504, 2014.
- [69] S. J. Gershman and Y. Niv, “Novelty and inductive generalization in human reinforcement learning,” *Topics in Cognitive Science*, vol. 7, no. 3, pp. 391–415, 2015.
- [70] A. Gopnik, “Childhood as a solution to explore–exploit tensions,” *Philosophical Transactions of the Royal Society B*, vol. 375, no. 1803, p. 20190502, 2020.
- [71] D. Zhou, D. M. Lydon-Staley, P. Zurn, and D. S. Bassett, “The growth and form of knowledge networks by kinesthetic curiosity,” *arXiv preprint arXiv:2006.02949*, 2020.
- [72] J. Schmidhuber, “Driven by compression progress: A simple principle explains essential aspects of subjective beauty, novelty, surprise, interestingness, attention, curiosity, creativity, art, science, music, jokes,” in *Workshop on anticipatory behavior in adaptive learning systems*, pp. 48–76, Springer, 2008.
- [73] A. E. Sizemore, E. A. Karuza, C. Giusti, and D. S. Bassett, “Knowledge gaps in the early growth of semantic networks,” *Nature Human Behavior*, vol. 2, no. 9, pp. 682–692, 2018.
- [74] H. Ju, D. Zhou, A. S. Blevins, D. M. Lydon-Staley, J. Kaplan, J. R. Tuma, and D. S. Bassett, “The network structure of scientific revolutions,” *arXiv preprint arXiv:2010.08381*, 2020.
- [75] C. W. Lynn and D. S. Bassett, “Compressibility of complex networks,” *arXiv preprint arXiv:2011.08994*, 2020.
- [76] P. Zurn and D. S. Bassett, “Network architectures supporting learnability,” *Philosophical Transactions of the Royal Society B*, vol. 375, no. 1796, p. 20190323, 2020.
- [77] A. Betti and H. van den Berg, “Modelling the history of ideas,” *British Journal for the History of Philosophy*, vol. 22, no. 4, pp. 812–835, 2014.
- [78] J. M. McCarthy and G. S. Soh, *Geometric Design of Linkages*. Springer-Verlag New York, 2011.
- [79] M. M, “A novel classification of planar four-bar linkages and its application to the mechanical analysis of animal systems,” *Phil. Trans. R. Soc. Lond.*, p. B351689–720, 1996.

- [80] J. C. Maxwell, “On reciprocal figures and diagrams of forces,” *Philosophical Magazine, 4th Series*, vol. 27, pp. 250–261, 1864.
- [81] J. Howard, *Mechanics of motor proteins and the cytoskeleton*. Sinauer Associates, 2001.
- [82] J. Z. Kim, Z. Lu, S. H. Strogatz, and D. S. Bassett, “Conformational control of mechanical networks,” *Nat. Phys.*, vol. 15, p. 714–720, 2019.
- [83] J. Z. Kim, Z. Lu, and D. S. Bassett, “Design of large sequential conformational change in mechanical networks,” *arXiv*, vol. 1906, p. 08400, 2019.
- [84] J. T. Overvelde, T. A. de Jong, Y. Shevchenko, S. A. Becerra, G. M. Whitesides, J. C. Weaver, C. Hoberman, and K. Bertoldi, “A three-dimensional actuated origami-inspired transformable meta-material with multiple degrees of freedom,” *Nat Commun*, vol. 10929, no. 7, 2016.
- [85] J. W. Rocks, N. Pashine, I. Bischofberger, C. P. Goodrich, A. J. Liu, and S. R. Nagel, “Designing allostery-inspired response in mechanical networks,” *Proc Natl Acad Sci*, vol. 114, no. 10, pp. 2520–2525, 2017.
- [86] C. R. Calladine, “Buckminster Fuller’s tensegrity structures and Clerk Maxwell’s rules for the construction of stiff frames,” *Int J Solids Struct*, vol. 14, p. 161–172, 1978.
- [87] D. Zhou, C. W. Lynn, Z. Cui, R. Ciric, G. L. Baum, T. M. Moore, D. R. Roalf, J. A. Detre, R. C. Gur, R. E. Gur, *et al.*, “Efficient coding in the economics of human brain connectomics,” *arXiv preprint arXiv:2001.05078*, 2020.
- [88] A. O. Constantinescu, J. X. O’Reilly, and T. E. J. Behrens, “Organizing conceptual knowledge in humans with a gridlike code,” *Science*, vol. 352, no. 6292, pp. 1464–1468, 2016.
- [89] M. M. Garvert, R. J. Dolan, and T. E. Behrens, “A map of abstract relational knowledge in the human hippocampal–entorhinal cortex,” *Elife*, vol. 6, p. e17086, 2017.
- [90] S. A. Ocko, K. Hardcastle, L. M. Giocomo, and S. Ganguli, “Emergent elasticity in the neural code for space,” *Proceedings of the National Academy of Sciences*, vol. 115, no. 50, pp. E11798–E11806, 2018.
- [91] S. H. Thompson, E. B. Falk, M. B. O’Donnell, C. N. Cascio, J. B. Bayer, J. M. Vettel, and D. S. Bassett, “Response inhibition in adolescents is moderated by brain connectivity and social network structure,” *Social Cognitive and Affective Neuroscience*, vol. 15, no. 8, pp. 827–837, 2020.
- [92] E. A. Karuza, A. E. Kahn, and D. S. Bassett, “Human sensitivity to community structure is robust to topological variation,” *Complexity*, vol. 2019, no. 8379321.
- [93] D. S. Bassett, P. Zurn, and J. I. Gold, “On the nature and use of models in network neuroscience,” *Nature Reviews Neuroscience*, vol. 19, no. 9, pp. 566–578, 2018.
- [94] E. A. Karuza, A. E. Kahn, S. L. Thompson-Schill, and D. S. Bassett, “Process reveals structure: How a network is traversed mediates expectations about its architecture,” *Scientific reports*, vol. 7, no. 1, pp. 1–9, 2017.
- [95] H. Ju and D. S. Bassett, “Dynamic representations in networked neural systems,” *Nature Neuroscience*, vol. 23, no. 8, pp. 908–917, 2020.
- [96] L. E. Suárez, R. D. Markello, R. F. Betzel, and B. Misic, “Linking structure and function in macroscale brain networks,” *Trends in Cognitive Sciences*, vol. 24, no. 4, pp. 302–315, 2020.
- [97] J. M. Huntenburg, P.-L. Bazin, and D. S. Margulies, “Large-scale gradients in human cortical organization,” *Trends in cognitive sciences*, vol. 22, no. 1, pp. 21–31, 2018.

- [98] D. S. Margulies, S. S. Ghosh, A. Goulas, M. Falkiewicz, J. M. Huntenburg, G. Langs, G. Bezgin, S. B. Eickhoff, F. X. Castellanos, M. Petrides, *et al.*, “Situating the default-mode network along a principal gradient of macroscale cortical organization,” *Proceedings of the National Academy of Sciences*, vol. 113, no. 44, pp. 12574–12579, 2016.
- [99] M. J. Gruber, A. Valji, and C. Ranganath, “Curiosity and learning: a neuroscientific perspective,” 2019.
- [100] C. Kemp and J. B. Tenenbaum, “The discovery of structural form,” *Proceedings of the National Academy of Sciences*, vol. 105, no. 31, pp. 10687–10692, 2008.
- [101] K. L. Stachenfeld, M. M. Botvinick, and S. J. Gershman, “The hippocampus as a predictive map,” *Nature neuroscience*, vol. 20, no. 11, p. 1643, 2017.
- [102] A. C. Schapiro, N. B. Turk-Browne, M. M. Botvinick, and K. A. Norman, “Complementary learning systems within the hippocampus: a neural network modelling approach to reconciling episodic memory with statistical learning,” *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 372, no. 1711, p. 20160049, 2017.
- [103] N. H. Christianson, A. S. Blevins, and D. S. Bassett, “Architecture and evolution of semantic networks in mathematics texts,” *Philosophical Transactions of the Royal Society, A*, vol. 476, no. 2239, p. 20190741, 2020.
- [104] M. W. Rossiter, “The Matthew/Matilda Effect in science,” *Social Studies of Science*, vol. 23, no. 2, pp. 325–341, 1993.
- [105] S. Ahmed, *On Being Included: Racism and Diversity in Institutional Life*. Duke University Press, 2012.
- [106] C. Mott and D. Cockayne, “Citation matters: mobilizing the politics of citation toward a practice of ‘conscientious engagement’,” *Gender, Place and Culture*, vol. 24, no. 7, pp. 954–973, 2017.
- [107] M. Fricker, *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford University Press, 2009.
- [108] N. Caplar, S. Tacchella, and S. Birrer, “Quantitative evaluation of gender bias in astronomical publications from citation counts,” *Nature Astronomy*, vol. 1, p. 0141, June 2017.
- [109] J. D. Dworkin, K. A. Linn, E. G. Teich, P. Zurn, R. T. Shinohara, and D. S. Bassett, “The extent and drivers of gender imbalance in neuroscience reference lists,” *Nature Neuroscience*, vol. 23, pp. 918–926, 2020.
- [110] D. Maliniak, R. Powers, and B. F. Walter, “The Gender Citation Gap in International Relations,” *International Organization*, vol. 67, pp. 889–922, Oct. 2013.
- [111] M. L. Dion, J. L. Sumner, and S. M. Mitchell, “Gendered Citation Patterns across Political Science and Social Science Methodology Fields,” *Political Analysis*, vol. 26, pp. 312–327, July 2018.
- [112] M. A. Bertolero, J. D. Dworkin, S. U. David, C. L. Lloreda, P. Srivastava, J. Stiso, D. Zhou, K. Dzirasa, D. A. Fair, A. N. Kaczkurkin, B. J. Marlin, D. Shohamy, L. Q. Uddin, P. Zurn, and D. S. Bassett, “Racial and ethnic imbalance in neuroscience reference lists and intersections with gender,” *bioRxiv*, vol. 10.12.336230, 2020.
- [113] S. M. Mitchell, S. Lange, and H. Brus, “Gendered Citation Patterns in International Relations Journals1,” *International Studies Perspectives*, vol. 14, no. 4, pp. 485–492, 2013.
- [114] P. Zurn and R. N. C. Bassett, Danielle S., “The citation diversity statement: A practice of transparency, a way of life,” *Trends Cogn Sci*, vol. 24, no. 9, pp. 669–672, 2020.

- [115] J. Dworkin, P. Zurn, and D. S. Bassett, “(in)citing action to realize an equitable future,” *Neuron*, vol. 106, no. 6, pp. 890–894, 2020.
- [116] D. Zhou, E. J. Cornblath, J. Stiso, E. G. Teich, J. D. Dworkin, A. S. Blevins, and D. S. Bassett, “Gender diversity statement and code notebook v1.0,” 2020.
- [117] A. Ambekar, C. Ward, J. Mohammed, S. Male, and S. Skiena, “Name-ethnicity classification from open sources,” in *Proceedings of the 15th ACM SIGKDD international conference on Knowledge Discovery and Data Mining*, pp. 49–58, 2009.
- [118] G. Sood and S. Laohaprapanon, “Predicting race and ethnicity from the sequence of characters in a name,” *arXiv preprint arXiv:1805.02109*, 2018.