

Architectural styles of curiosity in global Wikipedia mobile app readership

Dale Zhou^{1,2}, Shubhankar P. Patankar¹, David M.
Lydon-Staley^{1,4}, Perry Zurn⁵, Martin Gerlach^{3,10,11} and Dani
S. Bassett^{1,2,6,7,8,9,10,11}

¹Department of Bioengineering, University of Pennsylvania,
Philadelphia, 19104, PA, USA.

²Neuroscience Graduate Group, Perelman School of Medicine,
University of Pennsylvania, Philadelphia, 19104, PA, USA.

³Wikimedia Foundation, San Francisco, CA, USA.

⁴Annenberg School of Communication, University of
Pennsylvania, Philadelphia, 19104, PA, USA.

⁵Department of Philosophy, American University, Washington,
20016, DC, USA.

⁶Department of Electrical & Systems Engineering, University of
Pennsylvania, Philadelphia, 19104, PA, USA.

⁷Department of Neurology, University of Pennsylvania,
Philadelphia, 19104, PA, USA.

⁸Department of Psychiatry, University of Pennsylvania,
Philadelphia, 19104, PA, USA.

⁹Santa Fe Institute, Santa Fe, 87501, NM, USA.

¹⁰Corresponding authors: mgerlach@wikimedia.org and
dsb@seas.upenn.edu.

¹¹These authors contributed equally.

Abstract

Intrinsically motivated information seeking is an expression of curiosity believed to be central to human nature. However, most curiosity research relies on small, Western convenience samples. Here, we expand an analysis of a laboratory study with 149 participants browsing Wikipedia to 482,760 readers using Wikipedia's mobile app in 14 languages from

50 countries or territories. By measuring the structure of knowledge networks constructed by readers weaving a thread through articles in Wikipedia, we provide the first replication of two distinctive architectural styles of curiosity: that of the busybody and of the hunter. Further, we find that latent dimensions of browsing patterns explain the variation in knowledge network structure, and reflect the existence of a third architectural style of curiosity—the dancer—which was previously predicted by a historico-philosophical examination of texts written across the last 2 millennia and is characterized by creative modes of knowledge production. Finally, across languages and countries, we identify novel associations between the structure of knowledge networks and population-level indicators of spatial navigation, education, mood, well-being, and inequality. Collectively, these results advance our understanding of Wikipedia’s readership globally and demonstrate how cultural and geographical properties of the digital environment relate to different styles of curiosity.

Keywords: curiosity, network science, natural language processing, knowledge networks, Wikipedia

1 Introduction

As epitomized by a view of our species as *informativores* [1], humans need to forage for information just like omnivores need to forage for food [2, 3]. This need is captured by the notion of curiosity [4–7], a multi-faceted psychological construct characterized by a range of behaviors. Curiosity is an intrinsic motivation to seek, sample, and search for novel, uncertain, and complex information [8–10]. In an effort to capture the rich dynamics and functions of curiosity, its practice has been characterized by three architectural styles—the busybody, the hunter, and the dancer—excavated from texts written over the last two millennia using a historico-philosophical method [11]. The busybody scouts for loose threads of novelty, the hunter pursues specific answers in a projectile path, and the dancer leaps in creative breaks with tradition across typically siloed areas of knowledge [12]. These three architectural styles underscore a dimensional approach to the study of curiosity, foregrounding the practice of curiosity as an individual difference.

Despite marked progress in the field, the generalizability of curiosity research is limited due to modest samples of convenience that are not representative of the variation across diverse cultures. This limitation represents a general problem of over-reliance on studies using the English language [13], as well as Western, Educated, Industrialized, Rich, and Democratic (WEIRD) populations [14]. Indeed the historico-philosophical approach that uncovered the patterns of the hunter, busybody, and dancer predominantly surveyed the Western canon [15]. In contrast, curiosity itself can be an equalizing force for the knowledge and experiences that are often neglected or oppressed [16, 17].

It can support justice and equality by unveiling the *status quo* and inventing ideas for deconstructing and rebuilding current structures [18].

A focus on curious practice is crucial to detail curiosity's role in facilitating health and well-being, a role that has become increasingly complex as more nuanced perspectives on curiosity have emerged. High trait curiosity is viewed as a character strength due to work showing that curiosity supports learning and well-being [19]. For example, in an educational context, curiosity constitutes a strong motivator for memory, learning, and the creation of knowledge [20–25]. There is also evidence that, within the online ecosystem, curious individuals are better able to critically assess the novelty and quality of false information [26], to read deeply [27], and to explore others' perspectives, thoughts, and feelings [28]. More generally, curiosity has been associated with individual well-being, lower depressed mood, and greater openness to experiences [10, 29–31]. However, specific styles of curiosity may also be associated with depression, anxiety, and attention-deficit/hyperactivity disorder [32, 33]. Moreover, curiosity can be double-edged, with high curiosity associated with behaving erroneously yet overconfidently, being less discerning of information quality, lacking intellectual humility [26], and having more addictive behavior [34]. Taking a multidimensional approach to curiosity provides insight into the complex association between curiosity and well-being. For example, recent work demonstrates that the association between curiosity and well-being depends on the type of curiosity under consideration. Joyous exploration is a facet of curiosity associated with pure enjoyment of novel epistemic stimuli [10], which is consistently positively associated with indicators of well-being. By contrast, no such association emerges between deprivation sensitivity—a facet of curiosity associated with a motivation to seek information to overcome the feeling of being deprived of knowledge [35]—and well-being [36]. Thus, investigating the advantages and disadvantages of different *styles* of curious practice is key to understanding human behavior writ large [33].

To study the architecture, dynamics, and benefits of curiosity across diverse cultures and from the perspectives of information foraging theories, we operationalize the practice of curiosity using the observational perspective of knowledge networks [15, 37]. In this framework, the behavioral expression of curiosity is characterized by the construction of knowledge networks by individual readers [38]. We define the network's nodes as articles that readers access, and we define edges as the presence or absence of hyperlinks between articles. Individual differences in knowledge network building are assessed by measuring topological indicators of the three architectural styles: hunter, busybody, and dancer. Hunters build tight, constrained networks whereas busybodies build loose, broad networks; precisely how to characterize the networks of the dancer has remained elusive. The advantage of the network approach is that it can be used to develop models [39, 40] and to formalize classical psychological and historico-philosophical taxonomies of curiosity [11].

Recently, these methods were successfully applied in a laboratory study wherein 149 participants were asked to browse Wikipedia for 15 minutes a

day for 21 days [41]. Participants implicitly built different architectures of knowledge networks as they wove a temporal thread through visited articles. Hunters, in contrast to busybodies, build tighter and denser knowledge networks associated with their deprivation sensitivity, an aversive state of curiosity that motivates one to eliminate gaps in knowledge [10, 35, 42].

Here, we generalize and expand the analysis of knowledge network structure by building a novel dataset consisting of 482,760 Wikipedia readers accessing 14 different language editions from 50 countries or territories (**Figure 1A**). Wikipedia provides an ideal platform for the systematic study of curiosity because it provides large, representative samples of naturalistic information seeking. It is the largest encyclopedia, publishing over 60 million free articles in more than 300 languages [43]. Wikipedia is the most popular website (alongside search engines, Facebook, and YouTube) in 43 countries, outpacing any other website [44], attracting 1.5 billion unique devices and billions of pageloads every month [43]. People turn to Wikipedia when they need to learn about novel, uncertain, and complex information [45, 46]. As a result, Wikipedia is not only an encyclopedia that anyone can share in but also a record of an epochal shift in how humans digitally collect, organize, store, and disseminate information. In addition, Wikipedia has produced a digital environment which is one of the most important laboratories for social scientific research in history [47].

Using these data, we make three main contributions. First, we replicate the identification of hunter and busybody styles of curiosity by generalizing the framework of knowledge networks to a large, naturalistic population of Wikipedia’s readership across diverse cultures. Second, we systematically examine the variability of metrics characterizing knowledge-network building across countries, languages, and topics. In so doing, we show that building broad and diverse networks is—on a population level—predictive of spatial navigation, well-being, and both gender and education equality. Third, we expand the empirical characterization of knowledge networks by providing new quantitative evidence for the creative architectural style of the dancer. In making these three contributions, we provide new insights into how readers seek information on Wikipedia, we fundamentally expand the state-of-the-art in how curiosity practice is studied, and we deepen our understanding of human curiosity across cultures and languages.

2 Results

Generalizing knowledge networks

We first aim to generalize findings about curiosity inferred from knowledge networks of readers. To do so, we compare the knowledge networks from data collected in the laboratory at the University of Pennsylvania to the naturalistic knowledge networks of readers of Wikipedia across many languages and countries (see **Methods** for details). In the first dataset from UPenn, 149 participants browsed 18,654 unique Wikipedia articles. Participants were

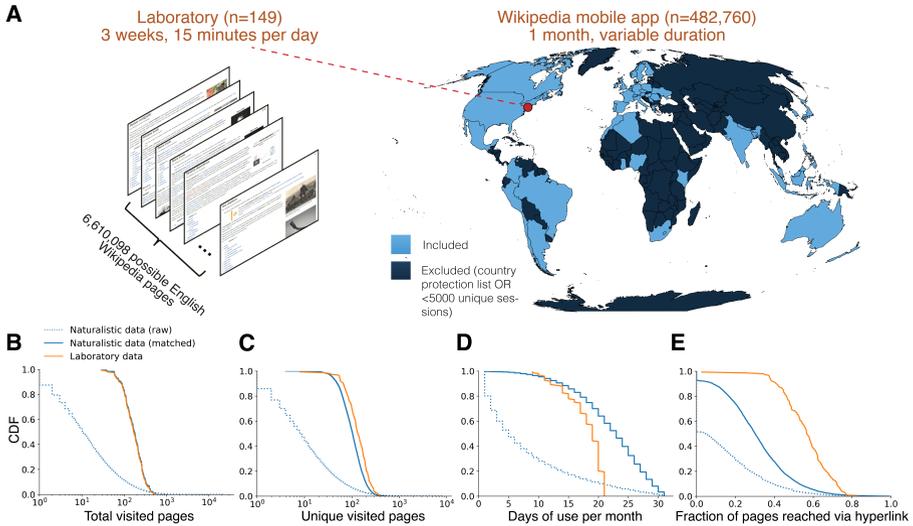


Fig. 1 Geography and summary statistics of laboratory versus naturalistic data. (A) Geography and sample sizes of laboratory and naturalistic data. Naturalistic data includes 14 languages and 50 countries or territories. (B–D) Cumulative distribution functions (CDF) of summary statistics for laboratory data and naturalistic data from English Wikipedia. To fairly compare knowledge networks between the datasets, we used propensity score matching, a quasi-experimental method that attempts to reduce the bias introduced by confounding variables, on the full dataset (dashed blue) to produce a biased sample with more similar summary statistics (solid blue) to the laboratory data (orange).

instructed to browse English Wikipedia for 15 minutes each day over a period of 3 weeks. In the naturalistic dataset, 2,111,851 readers browsed 16,385,207 articles through the Wikipedia mobile app (**Figure 1A**). On average, readers had 30.93 ± 87.39 unique article views over 7.97 ± 7.75 days. This second, naturalistic dataset was purely observational without any instructions to readers. To assess similarities and differences between the two datasets, we consider four summary statistics (**Figure 1**): the number of articles visited, the number of unique articles visited, the number of days of usage in 1 month, and the fraction of articles reached via Wikipedia hyperlinks (versus from an external website). We find several notable differences: on average, readers in the naturalistic dataset visited fewer articles, visited fewer unique articles, visited articles on fewer days, and reached a lower fraction of articles through internal hyperlinks (**Figure 1B–D**).

Given these differences in summary statistics, we select a subsample of 482,760 readers from the naturalistic dataset matched for page views to the laboratory dataset in order to fairly compare knowledge networks. Specifically, we use propensity score matching on the number of page views per reader (see **Methods** for details). Propensity score matching is a quasi-experimental method that attempts to reduce the bias introduced by confounding variables by creating an artificial control group from observational data [48]. While this procedure leads to a smaller subsample of readers, interpreting similarities

and differences between the two datasets becomes more feasible, based on network characteristics that are no longer driven by differences in the number of pages viewed per reader. As a result of propensity score matching, the distributions of the number of visited pages per reader are almost identical (**Figure 1B-D**). Furthermore, the differences in the distributions of the other summary statistics also become less pronounced, suggesting that differences in the datasets prior to propensity score matching derive largely from differences in the number of pages viewed per reader. Henceforth, we consider the matched subsamples from propensity score matching unless otherwise stated.

Assessing topological similarities

Next, we test how knowledge networks in the laboratory data generalize to the naturalistic data by systematically comparing their topological properties. In previous studies, curiosity in knowledge networks was characterized by the clustering coefficient and the characteristic path length [41]. These metrics measure architectural styles of curiosity that generate tight or loose knowledge networks. Here, we expand this characterization by calculating 6 additional network metrics (see **Methods** for details): degree, global efficiency, core-periphery structure, modularity, number of groups (or modules), and minimum description length. In considering the marginal distributions of each metric, we find that the two datasets are qualitatively similar (**Figure 2**). Particularly notable is the modularity metric, for which both distributions show a bimodal structure that peaks at approximately the same scales. This quantitative assessment of topological structure provides evidence that curiosity styles which have been uncovered in previous studies using the laboratory data generalize to a broader population of Wikipedia’s readership in the naturalistic data.

In order to quantify the similarity of the population of knowledge networks in the laboratory and naturalistic datasets, we calculate the distance $d \in [0, 1]$ between the respective distributions based on the commonly-used Kolmogorov-Smirnov test (see **Methods** for details). We find that the averaged distance between laboratory and naturalistic data is around $d \approx 0.3$ (**Figure 3A**). To interpret this value, we perform a similar comparison with reference knowledge networks generated from another dataset: Wikispeedia, an experiment where participants performed targeted navigation from a random source page to a target page [49, 50]. We find that the averaged distance between the naturalistic data and Wikispeedia is $d \geq 0.7$, significantly larger than the distance between naturalistic data and laboratory data. Next, we perform a similar comparison with reference knowledge networks generated from Wikipedia 6 months prior. We find that the averaged distance between the two naturalistic datasets was $d \leq 0.05$, significantly smaller than the distance between the naturalistic and laboratory data. Collectively, these results indicate that the naturalistic data are more similar to the laboratory data than they are to a targeted navigation task, and less similar to the laboratory data than they are to naturalistic data acquired at a different time point.

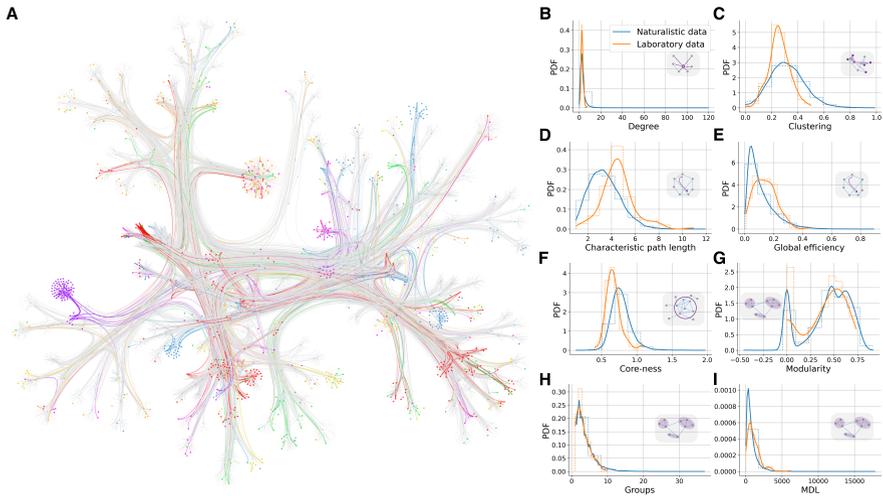


Fig. 2 Distributions of network metrics in knowledge networks from the laboratory and naturalistic data. **(A)** As readers seek information, they weave diverse temporal threads through nodes defined as Wikipedia articles and across edges defined as hyperlinks. Displayed is a hyperlink network with only 0.1% of nodes in English Wikipedia displayed and different knowledge networks of 7 readers highlighted in red, orange, yellow, green, blue, indigo, and violet. Network organized by hierarchical block partition, positioned using force directed placement, and with edges subsampled for visualization. **(B–I)** Solid lines indicate the probability density functions (PDF) from a kernel density estimation, whereas dotted lines indicate normalized histograms. Blue: naturalistic data. Orange: laboratory data. Insets provide conceptual depictions of the network metric. Metrics include **(B)** degree, **(C)** clustering, **(D)** characteristic path length, **(E)** global efficiency, **(F)** core-ness, **(G)** modularity, **(H)** groups, or number of modules, and **(I)** minimum description length.

Nevertheless, knowledge networks in the naturalistic data show a larger variation when stratified by the country from where readers access the English version of Wikipedia (**Figure 3B**). Although for most countries (such as Canada, the United States of America, Australia, and Great Britain) the distances are similarly small ($d \approx 0.05$), for other countries—such as Germany—distances are large and comparable to the distance between the laboratory and naturalistic data ($d \approx 0.2$). To account for the fact that readers from countries such as Germany will preferentially access the German version of Wikipedia (instead of the English version), we create naturalistic datasets for 14 language versions of Wikipedia (see **Methods**). Comparing the knowledge networks for the naturalistic datasets from different languages (**Figure 3C**), we find that distances between English and most languages (German, Russian, Japanese, Spanish, Dutch, Hebrew, Ukrainian, Chinese) are of comparable magnitude to the distance between the laboratory and naturalistic data ($0.1 \lesssim d \lesssim 0.3$). Some languages (Arabic, Hungarian, Bengali, Romanian, Hindi) show slightly larger distances ($0.3 \lesssim d \lesssim 0.4$). However, the latter can be partly explained by the fact that the underlying networks of hyperlinks between articles are much smaller in size (see **Supplementary Materials**).

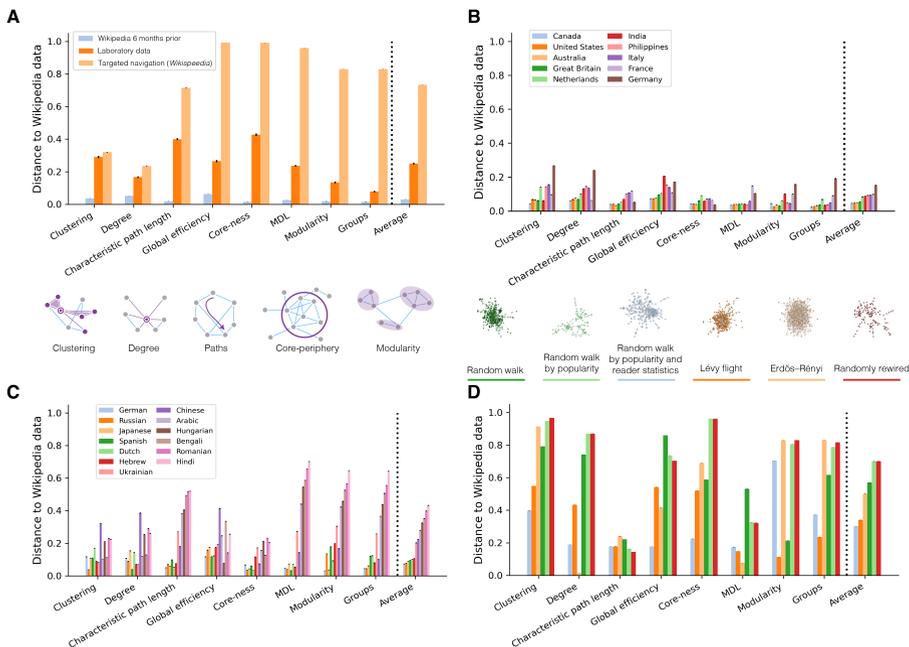


Fig. 3 Distances of the distribution of knowledge network structures between the naturalistic dataset and other datasets. **(A)** Comparison with three other datasets: a dataset of reading sessions from 2022-03 (instead of 2022-10), a lab-based Wikipedia navigation game called *Wikispeedia*, and laboratory data. **(B)** Knowledge networks from readers of English Wikipedia in specific countries. **(C)** Knowledge networks from readers of Wikipedia in other languages. **(D)** Different null model datasets synthetically generated from random walks or from network formation processes. The mean Kolmogorov-Smirnov distance is displayed with two-tailed 95% bootstrap confidence intervals from 100 iterations.

Comparing to null models

To assess the degree to which similar knowledge networks could arise by chance, we compare the naturalistic data with knowledge networks from different null models (see **Methods**). Random rewirings of empirical networks were used to create a benchmark for architectures that appear simply by chance. Erdős-Rényi random networks with matched edge density were used as benchmarks for architectures agnostic to any particular topological properties. Barabási-Albert networks—generated by a preferential attachment (rich-get-richer) mechanism—were used as benchmarks for architectures with heterogeneous degree distributions. We find that distances to these three null models ($d \approx 0.8$) are much larger than the distances to the laboratory data (**Figure 3D**). This result suggests that the similarity between the laboratory and naturalistic datasets is not merely a coincidence, nor easily explained by a rich-get-richer mechanism.

To assess the degree to which similar knowledge networks could result from different generative mechanisms, we compare the naturalistic data with knowledge networks generated from four types of random walks. Random walks have

previously been shown to resemble the behavior of hypothesized models and thus serve as important null models [51]. The first random walk is biased for the popularity of pages (as measured by cumulative page views) and takes into account the fact that 37.5% of page-to-page transitions (or navigations) were made by hyperlinks rather than by external sources ($d \approx 0.3$) [52]. The second random walk is biased to undergo an efficient foraging search called a Lévy flight modeled after the laboratory data ($d \approx 0.4$) [41]. The third and fourth random walks only navigated by random hyperlinks or only by the popularity of pages led to by the hyperlinks ($d \approx 0.7$). Collectively, these comparisons suggest that possible generative mechanisms for Wikipedia browsing in the naturalistic setting include Lévy flight dynamics and a random walk biased for popularity and source.

The performance of the Lévy flight model was of particular note for two reasons. First, in contrast to the popularity- and source-biased random walks, the Lévy flight model was fit to the laboratory data and then applied to the naturalistic data. As such, the degree of model fit is striking evidence of generalization across the laboratory and naturalistic contexts. Second, Lévy flight dynamics have rich cognitive and affective associations both theoretically and in the context of the laboratory data. Theoretically, in Lévy flight foraging, short steps serve to build closely connected clusters while long leaps connect clusters to novel spaces. Rather than foraging for physical resources, Lévy flights here may support foraging for information and psychological resources. Consistent with this idea, we find marked associations between Lévy flight parameters and the affective state of participants in the laboratory study: Individuals whose foraging was more constrained, due to being more repetitive, reported higher depressed mood ($\beta = 1.06$, $p = 0.02$) and anxiety ($\beta = 1.17$, $p = 0.02$) than individuals whose foraging was less constrained. These correlations between negative mood and constrained exploration are consistent with psychological theories that link negative mood to a narrowed or ruminative scope of attention [53–55]. They also raise the question of how knowledge networks built by Wikipedia readers relate to well-being, a question we return to later in this analysis.

Conservation of curiosity styles

Overall, the results thus far show that knowledge networks from the naturalistic data are very similar to those from the laboratory data reported in earlier studies. This similarity is stronger between the naturalistic and laboratory data compared to reference or null models (**Figure 3D**). This finding implies that curiosity styles uncovered from knowledge networks in the laboratory data can be generalized to the population at large in the naturalistic data. We now turn to a demonstration of this generalization by identifying the hunter and busybody styles of curiosity, characterized by distinct practices of building knowledge networks, in the naturalistic data [11, 41]. Specifically, we define a scalar busybody/hunter score by aggregating different network metrics; the score captures whether a knowledge network is more

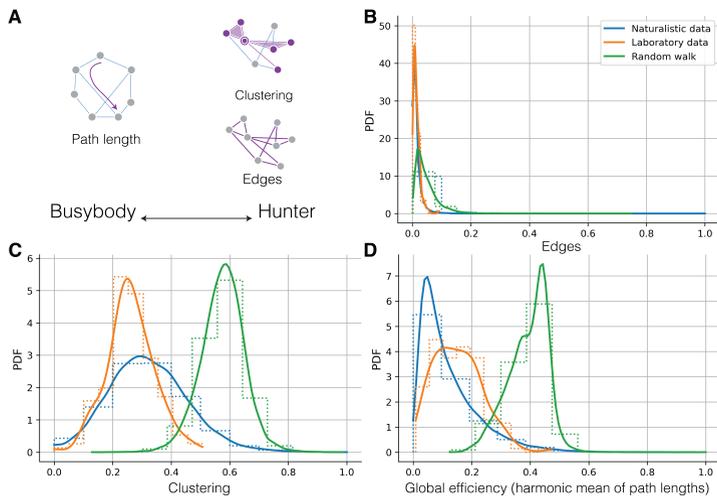


Fig. 4 Similar distributions of busybodies and hunters between the naturalistic and laboratory datasets. (A) The busybody/hunter score is a single metric that characterizes the clustering coefficient and characteristic path lengths in knowledge networks. These metrics were selected because they correlate with deprivation curiosity in the laboratory data [41]. The metric is the average number of edges, clustering coefficient, and global efficiency minus the characteristic path length. The naturalistic and laboratory data are more similar to each other than a random walk in terms of (B) number of edges, (C) clustering coefficient, and (D) global efficiency.

clustered and hunter-like, or more dispersed and busybody-like (Figure 4A; see Methods). We find that the laboratory and naturalistic data show similar distributions of the busybody/hunter score where d of edges, clustering, and global efficiency $\approx 0.18, 0.28$, and 0.26 respectively. These distances are markedly lower than those of the same measures for the laboratory and naturalistic datasets compared to the null networks generated by a random walk ($d_{\text{laboratory}} \approx 0.55, 0.95, 0.86$; $d_{\text{naturalistic}} \approx 0.59, 0.79, 0.86$; Figure 4B-D). The similar distributions of the busybody/hunter score in the naturalistic and laboratory datasets supports the generalizability of the hunter and busybody styles of curiosity.

Curiosity styles differ by topic

Readers visit Wikipedia for a variety of reasons and with different information needs [45, 46]. Might their curiosity styles differ according to those needs [56]? In this section, we seek to answer that question by exploring how readers' knowledge networks vary with respect to the topic that the reader is exploring. The individual needs of readers are reflected in what articles they choose to visit. For example, readers of English Wikipedia can choose from among more than 6 million articles on a wide range of topics. The reader's choice can be driven by epistemic curiosity, or a desire for new knowledge [4], which could mean exploring more of the same topic or exploring a diverse set of topics. Prior work has hypothesized that busybodies explore a richer diversity of topics

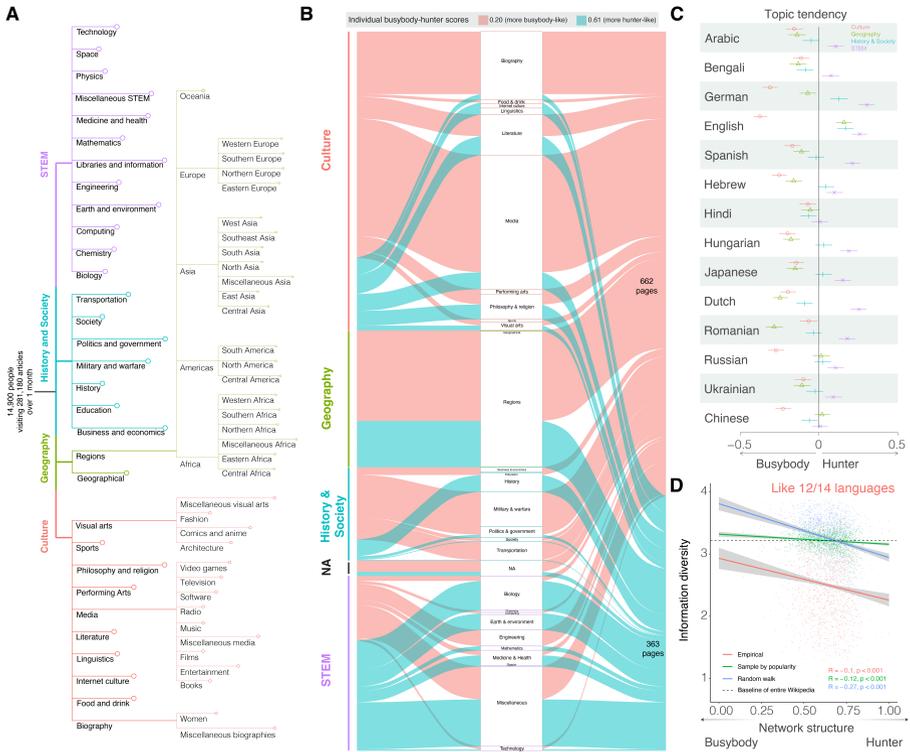


Fig. 5 Curiosity styles evinced by knowledge networks vary by topic. (A) The hierarchy of topics we considered, including 4 broad topics and 64 granular topics. (B) Alluvial diagram of page views for the 4 broad topics corresponding to the 64 granular topics for a user that was more busybody-like (salmon) or more hunter-like (cyan). (C) Forest plot of page views (normalized by the page views expected under a random walk) versus the busybody/hunter-score separately for the 4 broad topics. The mean Spearman's correlation coefficient is displayed with 95% bootstrap confidence intervals from 10,000 iterations. (D) Scatter plot of information diversity versus the busybody/hunter-score across the 64 granular topics. The empirical data is given by the red line. The popularity-biased random walk model is given by the green line. An unbiased random walk is given by the blue line.

than hunters [41]. Here we seek to test that hypothesis by characterizing how readers engage different topics as they browse Wikipedia.

We begin by considering Wikipedia's topic taxonomy which groups articles into a small set of pre-defined topics ([57]; see **Methods** for details). For example, high-level topics span four main areas: STEM, History & Society, Geography, and Culture (**Figure 5A**). By examining knowledge networks with relatively high and low values of the busybody/hunter score, we qualitatively observe different topical preferences in the corresponding articles (**Figure 5B**). To assess this trend quantitatively, we calculate the correlation between a given curiosity style (busybody/hunter score) and the number of articles being read that belong to each of the four broad topics. We normalize the number of articles being read from each topic by the number that would be expected

from a random walk to identify correlations that are significant and differ from a null model. Across 14 languages, we consistently find a relationship between the hunter and busybody styles and the number of articles being read from the four broad topics (**Figure 5C**). Readers who were more similar to busybodies tended to read articles about culture topics (media, food, art, philosophy, religion, etc.; 14/14 languages) and geography (11/14 languages), more so than expected in the null models. Further, readers who were more similar to hunters tended to read articles about STEM topics (12/14 languages), more so than expected in the null models. Hunters tended to read articles about History & Society in some languages (German and English), whereas busybodies tended to do so in other languages (Arabic, Bengali, Hindi, Dutch, and Chinese). These tendencies are consistent with the hypothesis that busybodies gravitate more towards social topics than do hunters [11].

In addition to the topic being engaged, it is interesting to assess whether the curiosity style of the reader relates to the diversity of topics being explored. We find that hunter-style knowledge networks display a lower diversity of topics than busybody-style knowledge networks (**Figure 5D**). This relationship is statistically evident from the significant correlation between the busybody/hunter score and the information diversity score (Spearman rank correlation coefficient: $\rho = -0.10$, $p < 0.001$), and this trend applies to 12 of the 14 languages we examined (see **Methods** for a definition of the information diversity score). Importantly, the diversity of topics browsed by humans is significantly lower than the diversity of topics canvassed by two null models; the random walk null model and the popularity-biased random walk null model.

Interestingly, in both null models, we observe a correlation between the busybody/hunter score and the information diversity score (Spearman rank correlation coefficient: $\rho = -0.12$, $p < 0.001$ for the popularity-biased random walk and $\rho = -0.27$, $p < 0.001$ for the random walk). This observation motivates the question of whether a feature of the hyperlink network could explain some variation in how the network is browsed. To address this question, we expand our assessment to different Wikipedia language editions. We find that one relatively consistent correlate of the busybody/hunter score is the size of the underlying hyperlink network between articles (see **Supplementary Materials**). This observation indicates that the structure of the underlying hyperlink network between articles can induce systematic differences in the observed curiosity style. However, we find that hyperlink network size alone does not fully explain the busybody/hunter score, as the correlations, for example between the number of culture articles visited and the busybody/hunter score exceeds that of null models matched for size. This observation demonstrates that some of the variance in observed curiosity style cannot be fully explained merely by differences in the underlying network; rather, readers with different curiosity styles also have preferential topical interests and browsing practices.

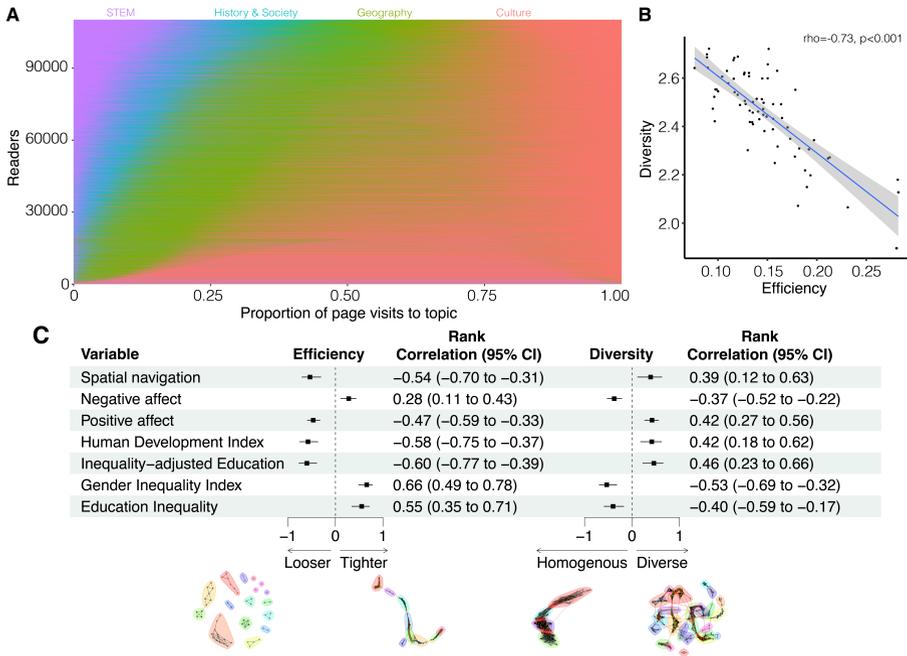


Fig. 6 Diversity and efficiency of knowledge networks predict population-level sociodemographic variables. (A) Readership across 74 pairs of countries and languages vary in the diversity of topics explored, where high diversity (top) corresponds to a more uniform representation of topics compared to low diversity (bottom) with an over-representation of some topics (4 broad topics visualized, analyses performed on 64 granular topics). (B) Networks built with more information diversity, defined as the average Shannon diversity index over 64 topics, tend to have reduced information efficiency, defined as the average global efficiency. (C) Looser networks have reduced global efficiency but are broader reaching. Looser networks were associated with better spatial navigation, mood, well-being, education, and equality. Diverse networks with greater Shannon diversity indices were associated with better spatial navigation, mood, well-being, education, and equality scores. The 95% bootstrap confidence interval is displayed from 10,000 iterations.

Network variation tracks population-level differences in spatial navigation, inequality, education, and well-being

In analyzing curiosity-driven architectural styles of knowledge network building, prior reports have found that such styles are predictive of sensation-seeking, an individual-level psychological state [41]. Our results thus far also suggested that repetitive and constrained styles of knowledge network building are predictive of individual-level states of negative affect. These observations motivate the question of whether knowledge networks built by Wikipedia readers track with population-level predictors of psychological states, particularly those related to well-being. Considering this question across geographical and cultural locales is important given evidence of the effect of one's environment and the environment's wealth and gender inequality on learning [58] and psychological assessments of spatial navigation, a universal requirement across

cultures [59]. To address this question, we examined whether and how tight, loose, or diverse knowledge networks across countries and languages relate to population-level indices of spatial navigation, education, life expectancy, and mood (see **Methods** for details). To account for the role of geography and cultural context, we consider measures averaged within datasets of 1,490 readers, each from 74 pairs of a given country and a given language edition of Wikipedia (**Figure 6A**). We examine population-level variables for spatial navigation, negative affect, and positive affect. We also examine population-level aggregate measures. We include the Human Development Index, which aggregates the life expectancy at birth, expected and average years of education, and income. We further consider the inequality in education, determined by the Atkinson index of the normalised ratio of the equally distributed level of education to the average observed distribution. Finally we consider gender inequality, which aggregates the maternal mortality ratio, adolescent birth rates, secondary education ratio, and labor force participation ratio. We test if these variables are related to the information diversity and the global efficiency of the network, measured as the inverse of the average shortest distance between all pages. We highlight this metric from the busybody/hunter score because a sensitivity analysis showed that the global efficiency of knowledge networks was most correlated with the information diversity of the network ($\rho = -0.73, p < 0.001$; **Figure 6B**).

We start by calculating average scores for global efficiency and information diversity across readers' knowledge networks for each of the 74 country-language pairs. Treating the average scores as independent variables and correcting 14 p -values with the Bonferroni method, we find that loose networks are associated with higher spatial navigation performance ($\rho = 0.54, p = 8.84 \times 10^{-6}$), less negative affect ($\rho = -0.28, p = 0.01$), more positive affect ($\rho = 0.47, p = 9.03 \times 10^{-8}$), greater Human Development Indices ($\rho = 0.58, p = 1.46 \times 10^{-6}$), more education ($\rho = 0.60, p = 3.25 \times 10^{-7}$), lower gender inequality ($\rho = -0.66, p = 7.46 \times 10^{-9}$), and lower education inequality ($\rho = -0.55, p = 6.74 \times 10^{-6}$; **Figure 6C**). In contrast, more diverse networks have no association with spatial navigation performance ($\rho = 0.39, p = 0.1$), but are associated with less negative affect ($\rho = -0.37, p = 7.57 \times 10^{-5}$), more positive affect ($\rho = 0.42, p = 3.31 \times 10^{-6}$), greater Human Development Indices ($\rho = 0.42, p = 3.86 \times 10^{-3}$), more education ($\rho = 0.46, p = 6.67 \times 10^{-4}$), lower gender inequality ($\rho = -0.53, p = 2.95 \times 10^{-5}$), and lower education inequality ($\rho = -0.40, p = 7.26 \times 10^{-3}$; **Figure 6C**). In general, the measure of network architecture (global efficiency) had stronger effect sizes of relationships with these dependent variables than a purely content-based measure of information diversity. These differences in effect sizes foreground the importance of modeling the complex connections coalescing from sequences of information beyond the content of that information. Taken together, the results are consistent with the notion that certain forms of curiosity support well-being [10, 60] and that policies and practices of equality may support less restricted and diverse forms of curiosity [17]. They also raise the question of whether the

social forces underlying inequality serve to constrain curiosity into hunter-like styles, thereby negatively impacting epistemic well-being.

Dancer type of curiosity

While we have thus far focused on the hunter and busybody styles of curiosity, a distinct signature of curiosity from the previous historico-philosophical examination that has yet to be operationalized is the dancer. This type of curiosity is described as a dance in which disparate concepts, typically conceived of as unrelated, are briefly linked in unique ways as the curious individual leaps and bounds across traditionally siloed areas of knowledge [11]. Such brief linking fosters the generation or creation of new experiences, ideas, and thoughts. Hence, it is of use to turn to the literature on creativity to seek computational frameworks in which to operationalize dancer-like curiosity. In psychology, generating creative ideas has previously been theorized to require a person to remodel traditional conceptual associations in favor of those associations that are not typically made [61]. Moreover, that remodeling process has been thought through the language of network science, raising the possibility that a network approach could offer a fruitful avenue by which to capture the curious practice of the dancer [62–64].

A recent methodological advance especially relevant to identifying the dancer signature is that of *forward flow* in information search. Forward flow characterizes the momentous “leaps” in thought that are predictive of creativity [62]. Forward flow is calculated by the average distance between the current “thought” and all previous thoughts. Prior work suggests that this metric may be useful in analyzing naturalistic text data; its application to social media posts can predict the scope of a poster’s creative achievements [62]. Here, we use forward flow to capture the creativity that is characteristic of the dancer type of curiosity as it connects seemingly disparate concepts.

As an illustration, we first examine a reader’s knowledge network and re-define edges as the semantic distance between articles (Figure 7A; see **Methods** for details). Consider two simulated walks weaving a temporal thread through 5 pages where one walk seeks information with high forward flow and the other with low forward flow; each follows qualitatively different trajectories and visits different types of pages (**Figure 7B**).

We then analyze the forward flow of Wikipedia readers. Because dancers must explore diverse topics to connect them, we examined the relation between forward flow and information diversity. We found that the two were significantly positively correlated ($\rho = 0.79$, $p < 0.001$). We also found that the distribution of forward flow in readers’ knowledge networks from English Wikipedia is different from that of null model knowledge networks generated from random walks ($d \approx 0.78$; **Figure 7C**), suggesting that empirical forward flow cannot simply be explained by random chance. Specifically, forward flow for readers is lower than for random walks, indicating that Wikipedia readers connect disparate topics in a more constrained way.

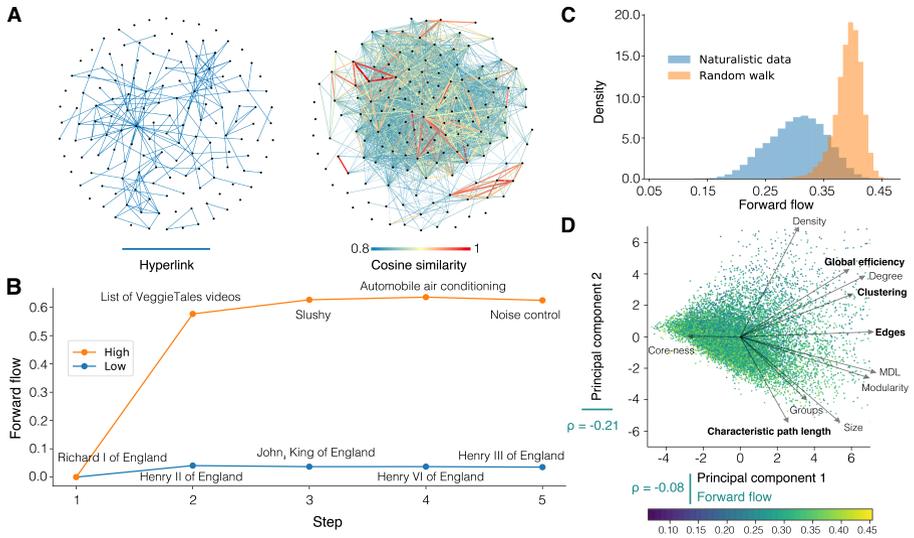


Fig. 7 Operationalizing the dancer style of curiosity in knowledge networks. (A) Forward flow is a topological measure associated with creative thought. The measure is based on a network where edges are defined as the semantic similarity between articles (right) rather than as hyperlinks (left). The semantic network is visualized with edge strengths thresholded above 0.8. **(B)** Example of the forward flow metric changing during 5 steps of a simulated walk with high forward flow and another with low forward flow. Forward flow is the average cosine dissimilarity between the current article and all previous articles. **(C)** Forward flow values of English Wikipedia readers are distributed differently than those of a random walk null model ($d \approx 0.78$). **(D)** Principal components analysis of 11 topological metrics across 14,900 English Wikipedia readers. Arrows depict the loadings that relate the 11 topological metrics to the principal components, with bolded metric names being used in the busybody/hunter score. Forward flow is weakly correlated with the most dominant direction of variance in network statistics (Spearman's $\rho = -0.08$, $p < 0.001$) and is more strongly correlated with the second most dominant direction of variance in network statistics than the busybody/hunter score ($\rho = -0.21$, $p < 0.001$). Importantly, forward flow is most weakly correlated with the busybody/hunter score ($\rho = -0.05$, $p < 0.001$). Hence, forward flow captures a different feature of curiosity.

We also observe that forward flow captures some of the topological variation of knowledge networks. Specifically, principal components analysis indicates that two dominant directions of variance explain 68% of variance in all 11 topological metrics (**Figure 7D**). The values of forward flow in knowledge networks correlate with the first two principal components in a dimensionality reduction of all 11 topological metrics, suggesting that hypothesized styles of curious practice are supported by latent dimensions underlying the growth and form of knowledge networks. Most importantly, forward flow is only weakly correlated with the busybody/hunter score ($\rho = -0.05$, $p < 0.001$). These results suggest that forward flow describes a distinct aspect of knowledge networks related to creativity, which goes beyond the known curiosity styles of the hunter and the busybody.

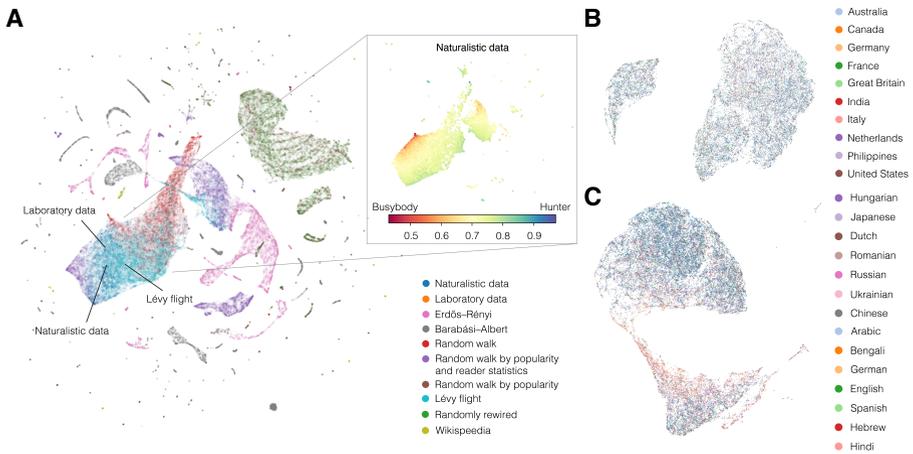


Fig. 8 Uncovering the space of curiosity styles in knowledge networks. (A) All 11 topological metric scores projected onto two dimensions using uniform manifold approximation and projection (UMAP) and colored by dataset or model. Naturalistic data for English Wikipedia is more proximal to the laboratory data and Lévy flight model than other null models (randomly rewired, random walks, Erdős-Rényi, Barabási-Albert) and browsing tasks (Wikipedia). Arrows label the centroids of the dataset. Inset: The naturalistic dataset from English Wikipedia is well-explained by the overlaid busybody/hunter score. (B-C) Different countries and languages tend to cluster together, suggesting comparable network structure. Two clusters separate due to network size differences (see **Supplementary Materials**).

Beyond the known styles of curiosity

Taking advantage of the large number of observations, we can uncover a rich spectrum of curiosity styles beyond the theoretically well-defined hunter, busybody, and dancer types [11]. We generate an unsupervised grouping (embedding) of the knowledge networks that shows the extent to which networks vary in their topological structure (Figure 8A). One main axis of this space is the distinction between hunter and busybody scores, from top left to bottom right (Figure 8A, inset). By including different null models, we can further understand how empirical knowledge networks are situated in the wider space of possible topologies. For example, some areas of the space are populated by networks built from random walks. Knowledge networks in English Wikipedia from different countries populate the same space of curiosity styles and are mixed homogeneously (Figure 8B). Knowledge networks from different Wikipedia languages populate similar spaces of curiosity styles. There are some systematic nuances between language versions (Figure 8C). Some of these differences can be explained by the underlying network structure, such as the different sizes of the hyperlink networks (see **Supplementary Materials**). This embedding approach quantifies the similarity of laboratory and naturalistic data, the dissimilarity of different models of curiosity, and the generalization across different countries and languages.

3 Discussion

Summary of findings

To summarize, we provide a systematic analysis of the architecture of knowledge networks constructed by readers of Wikipedia using the mobile app. We show that knowledge networks from previous small-scale lab-based studies are very similar in structure to those obtained from readers who are browsing Wikipedia in their daily lives. Critically, the distributions of network structures that characterize the hunter and busybody styles are similar between the laboratory and naturalistic data. The main distinction between the two datasets is that the latter not only constitutes a much larger (by orders of magnitude) and diverse sample (several countries and languages) but also that readers browsed Wikipedia without instructions by the authors of the study. The observed similarity in knowledge network structure between the two datasets is not a trivial finding: i) comparison with different null models demonstrates that it is unlikely to emerge by chance; and ii) previous studies have shown that navigation strategies of readers in Wikipedia can differ substantially from those observed in laboratory studies of targeted navigation such as the Wikispeedia game [52]. Our analysis also reveals new insights about the variation of knowledge networks in the Wikipedia readership. We show that knowledge networks display systematic differences in structure along different dimensions such as i) the topic space being explored; ii) the structure of the underlying link network; and iii) the cultural context of readers such as the country from where they access Wikipedia. Curiosity has been increasingly recognized to have a complex association with learning and well-being. Our operationalization of distinct curious practices captures this complexity and underscores the importance of doing so: differences in knowledge network structure and the diversity of topics explored are predictive of education, well-being, and equality. We further characterize knowledge networks beyond the known curiosity types of the hunter and the busybody: i) we operationalize the dancer type of curiosity by applying a forward flow measure previously associated with creativity; and ii) we use an unsupervised clustering technique to uncover a wide spectrum of curiosity styles that are as yet unexplored.

Implications

These findings have several important implications.

First, our work constitutes the first study replicating the network-based framework of curiosity. This work is therefore important in light of recent calls for prioritizing replication in research on human behaviour [65]. Our results allow us to generalize the framework of knowledge networks to a larger population outside of a laboratory setting. In turn, this replication enables us to expand the applicability of the characterization of curiosity derived from these networks. For example, we capture the continuum of styles along the

dimension that distinguishes the two paradigmatic types of curiosity denoted as the hunter and the busybody.

Second, our work provides new contributions to the theory of curiosity by revealing myriad outcome variables associated with different styles of curiosity. Historico-philosophical studies have identified three main types of curiosity: the hunter, busybody, and dancer [11]. We have proposed a new approach to quantify the latter in terms of momentous “leaps” that create a knowledge network. The connective movements in a conceptual space that spark creative insight may involve mechanisms of mental navigation much like that of spatial navigation [66, 67]. These mechanisms center around the mental maps we create, here operationalized in knowledge networks, by learning the network structure of a sequence of information [68–71]. Our results suggest that less fettered, more diverse, and broader information seeking—in contrast to the targeted wayfinding in games such as “Wikispeedia”—is related to decreased negative mood in both the laboratory and naturalistic data, as well as spatial navigation at the population level. Such broader and more diverse sequences are more consistent with the mind wandering and the creative thinking of a busybody and dancer than with the goal-directed and targeted thinking of a hunter [72]. These effects of enriched information seeking in digital environments are consistent with evidence that enriched environments can support some aspects of cognition and mental health [58]. We further tie the architectural styles of information seeking to efficient foraging theories [37, 73–76]. Different styles of curiosity are associated with different types of resources reaped, necessitating a reconfiguration of existing theories to account for a richer taxonomy of curiosity. For example, in the broaden-and-build model [55], positive emotions like interest tend to *broaden* exploration. Broadened exploration may enable people to *build* intellectual, emotional, and social resources [55, 77]. Curiosity may help one not only to build resources in the sense of accumulation but also to creatively draw on resources to reinvent and reconfigure the structure of their beliefs, self, and relationships with each other [40]. In contrast to broad openness, a need for closure may drive a tendency to “seize and freeze” information. The tendency to undiscerningly latch on to information that hastens the resolution of uncertainty, as opposed to protracted information gathering [78], may relate to inflexibly updating or making rigid the current structures of information [26, 79]. These theories suggest that there is wisdom in wandering in addition to the fruits of focused pursuit. Consistent with this suggestion, we observed that the practice of curiosity by people experiencing more negative emotions is narrowed; importantly, however, this is not evidence of narrow-mindedness. It is precisely the narrowed and returning types of movement that create more clustered networks associated with compressibility, navigability, and efficiency [80–82]. Such networks may be useful to regulate psychological uncertainty [41] and to help plan especially when the future environment is uncertain [76, 83]. The management of uncertainty [9, 84] and the constraints of returning to a central location (or home) versus exploring parts unknown [39, 82, 85] are core to theories of curiosity and how humans accumulate and

rely on shared information for innovation. Future work could incorporate the uncertainty of perception and memory as agents learn the structure of their environment [2, 76, 86]. Taken together, this framework provides a foundation for future data-driven studies on curiosity. The framework emphasizes that different styles of curiosity can have complex relationships to myriad outcomes including well-being and creativity depending on the context, richness, and uncertainty of physical and digital environments.

Third, our results have several practical implications. The framework of capturing readers' curiosity through their knowledge networks provides a more informative account of how readers engage with content on Wikipedia. Standard metrics of online engagement, such as session length [87], which are based on counting the number of pageviews can provide a misleading account of the information value that a website provides to readers [88]. In contrast to measuring only the quantity of engagement, operationalizing curiosity through knowledge networks can be used to measure *how* the readership engages with the content. Given recent studies that have reported associations between specific facets of curiosity and higher rates of errors in discerning the novelty and quality of information [26], such insights are relevant in addressing problem of dis- and misinformation in Wikipedia [89] from the perspective of readers. While most of the existing tools and initiatives have been geared towards supporting editors to maintain the quality of the content, we lack systematic approaches to support readers to become more resilient to unreliable content. However, Wikipedia is explicitly recommended as a valuable resource to consult in courses that teach how to evaluate the credibility of information online¹ using the technique of lateral reading across multiple sources of information [90]. Our results on the breadth and diversity of knowledge networks highlight the need to adapt similar strategies within the Wikipedia ecosystem. More generally, the framework of knowledge networks provides a starting point for the development of tools to support readers' exploration of Wikipedia given their different styles of curiosity. In turn, such support could lead to an improvement in serving readers' differing information needs and in helping them to easily find and access relevant information. In fact, items from the 5-dimensional curiosity (5DC) scale [7] overlap with the taxonomy of why readers visit Wikipedia [45]. For example, the 5DC-item measuring joyous exploration, "I seek out situations where it is likely that I will have to think in depth about something", matches one of the possible answers on the questionnaire about the information needs of readers in Wikipedia, "I am reading this article to get an in-depth understanding of the topic." Existing tools to recommend articles to readers in Wikipedia are based on popularity (top read articles shown in the Wikipedia app's feed²) or similarity (related articles shown at the bottom of the mobile site³). In contrast to these existing tools, knowledge networks provide a framework to formulate objective functions that could support learning. Recent work has shown how the construction

¹<https://cor.stanford.edu/videos/using-wikipedia/>

²https://www.mediawiki.org/wiki/Wikimedia_Apps/Feeds

³<https://www.mediawiki.org/wiki/Extension:RelatedArticles>

of knowledge networks is consistent with compression progress theory [40, 91]. In this view, curiosity is the drive to obtain new information to construct mental models with higher compression (i.e., lower cost of representation) which improves the capacity for abstraction and generalization believed to be crucial for human learning [39, 70]. Seeking information with more connections in learned mental models may support memory [92], perhaps because such connectivity increases the compressibility and navigability of networks [80, 81]. Future work could examine the effects of information seeking connectivity on well-being by analyzing longitudinal behavior and testing the effect of experiments that provide readers with page metrics about the information diversity and efficiency of their dynamic network structure. The results of such a study could inform (i) what we expect readers to learn as they search for information, (ii) how to improve the navigability of information networks, and (iii) how the connections and positively or negatively valenced content of recommended information is expected to affect memory and mood [52, 93–97]. On a population level, we found a statistically significant association between the global efficiency of knowledge networks—or more hunter- versus busybody-like patterns of Wikipedia browsing—and measures of well-being, education, and equality. This observation captures important differences in the relationships between the articles visited by readers. While we assessed the proportion of topics readers visited, future work could assess the transition probabilities between topics. Notably, the associations we observed were stronger for network efficiency than purely content-based measures such as information diversity, suggesting a need for new analyses of intensive longitudinal data of day-to-day digital experiences [98, 99]. Broadly, this analysis complements prior work that found relations between the motivations of Wikipedia readers and socio-demographic variables of the readers’ country such as the Human Development Index [46]. Our finding yields a deeper understanding of how a global readership consumes and connects knowledge across languages and countries. More importantly, the relation to socioeconomic variables highlights the necessity to take into account the local context of readers when developing initiatives to better serve their needs in alignment with the Wikimedia movement’s strategic effort towards knowledge equity⁴.

It is worthwhile to consider the associations between hunter-like curiosity and both gender and education inequality in light of prior work in critical social theory [100] and intersectional concerns of epistemic freedoms and knowledge production [101]. Candidate drivers of the curiosity-inequality relationship are many, and parsing their relative explanatory power would be an important goal for future work. Such drivers include gender, racial/ethnic, socioeconomic, and epistemic norms as well as their intersection. For example, it is possible that patriarchal forces [102, 103] drive a narrowing of curiosity practices—and a constraining of knowledge production approaches—that leads to hunter-like curiosity at the expense of other diverse forms of curiosity. A second possible explanation relates to gender disparities in readership: It is possible that in

⁴https://meta.wikimedia.org/wiki/Strategy/Wikimedia_movement/2017/Direction

countries characterized by greater gender and education inequality, the readers of Wikipedia tend to be people socialized as men, who are trained to evince hunter-like curiosity over busybody-like or dancer-like curiosity. A third possible explanation relates to epistemic norms in different disciplines or areas of knowledge. In particular, because of statistical associations between gender and discipline (whereby people socialized as women are pushed away from STEM), it is possible that countries with more men readers will evince styles of curiosity trained by STEM education, which may align more with the hunter-like style of curiosity than the busybody-like or dancer-like style. Across these candidate drivers, it would be interesting to test whether action to support diverse curiosity styles might effectively work against hegemonic norms. However, it is also important to acknowledge that readers in locations of marked social injustice (including gender and education inequality) could effectively utilize any sort of curiosity (including hunter-like curiosity) to build resistant knowledge [17, 104].

Lastly, our work demonstrates how Wikipedia is a human-centered project that continues to help advance our understanding of humans. Primarily, Wikipedia is an encyclopedia that anyone can share in. At the same time, it plays a crucial role in enabling research as “the most important laboratory for social scientific and computing research in history” [47]. Our work on the curiosity of Wikipedia’s readership constitutes an example of how this laboratory contributes to addressing current challenges in human behaviour research—not only by increasing ecological validity or diversifying representation empirically [105] but also by advancing theories of human behaviour [106]. More generally, Wikipedia plays a central role in contributing to the development and improvement of the scientific understanding of meaningful measures of human society [107]; as an encyclopedia and one of the largest global platforms for free knowledge it allows access to distinct aspects of human behavior complementing other widely-used resources that often observe behaviour on social media sites [108]. Naturalistic approaches may inform how future studies could strengthen the modest or incomplete relationships between performance on curiosity laboratory tasks and some facets of curiosity as a personality trait [109].

Limitations

Several methodological considerations and limitations are pertinent to this study. First, despite its size, the naturalistic dataset cannot be considered representative of all of Wikipedia’s readership. Here, we only consider readers using the Wikipedia app for whom we are able to construct knowledge networks that are sufficiently similar (e.g., same observational period) to the laboratory data. While most readers use the desktop or mobile-web platforms [52], our naturalistic mobile app dataset with millions of readers still constitutes the largest sample of knowledge networks currently available. Second, country-level aggregate measures like the Human Development Index are overly simplistic, neglect cultural differences, and have conceptual and methodological flaws

[110, 111]; for example, redundant sub-indices and false equivalences between indices of health, economic output, and education occur via averaging. Metrics of inequality address some but not all of these issues. Future work could consider expanding and replicating our results using additional data sets and aggregate measures.

4 Methods

Data

Laboratory

We used data from the Knowledge Networks Over Time (KNOT) study, a study designed to provide insight into behavior across a range of domains, including curiosity [30, 41, 60]. Participants ($n=149$; 121 women, 26 men, 2 genders beyond those listed as response options) were recruited through poster, Facebook, Craigslist, and university research site advertisements in Philadelphia and the surrounding university community. Participants were aged between 18.21 and 65.24 years (mean=25.05 years, s.d.=6.99 years) and identified as African American/Black (6.71%), Asian (25.50%), Hispanic/Latino (5.37 %), multiple races/ethnicities (5.37 %), races/ethnicities beyond those listed (5.37 %), white (49.66 %), and 2.01 % of participants did not report their races/ethnicities. Data collection began in October 2017 and ended in July 2018.

Interested participants were sent a baseline survey through Qualtrics that contained demographic and personality questionnaires. Participants then visited the laboratory and completed additional questionnaires, received training in a daily assessment protocol that began after the laboratory session, and were guided through the installation of tracking software (Timing) necessary for the Wikipedia browsing task which is the focus of the present manuscript.

After the laboratory visit, a 21-day daily assessment protocol was initiated. The 21-day assessment consisted of two components. Links to the daily assessments were emailed to participants at 18:30 each evening, and participants completed them outside of the laboratory on their personal computers.

The first component was a daily diary, delivered using Qualtrics, and consisted of survey questionnaires that took approximately five minutes to complete. These surveys included scales for depressed mood and anxiety [112].

The second component came immediately after the daily diary and was a 15-minute Wikipedia browsing task. As part of the Wikipedia browsing task, participants were prompted to open a browser and navigate to Wikipedia.org. The participants were instructed to spend 15 minutes in self-directed information seeking on Wikipedia and to explore whatever topics interested them. Specifically, during the laboratory visit, the investigator stated, "We would like you to open a new tab on your browser and visit <https://www.wikipedia.org/>. We would like you to spend 15 minutes each evening reading about whatever you want on Wikipedia. For example, if you wanted to learn more

about Philadelphia, you could go to the Philadelphia Wikipedia page.” At this point, the researcher used the Wikipedia search bar to navigate to <https://en.wikipedia.org/wiki/Philadelphia> to ensure that all participants had familiarity with Wikipedia and its usage. ”You can read through the page. You can also click on links you find interesting or you can use the search bar to search for new topics. There is no right or wrong way to do this. We are interested in what it is that people read about when they are not forced to read about anything in particular.” We developed this set of instructions to ensure that people would browse according to their curiosity and not in any particular manner suggested by the experimenter. After the 15 minutes of open browsing, the participants uploaded their browsing history.

Participants were compensated with Amazon gift cards at each study phase. They were given cards worth US\$25 after completing the baseline assessment and the laboratory visit. For the daily assessment, completion was incentivized by participant payments that were contingent upon completion: the completion of three, four, five, six, and seven surveys each week was compensated with gift cards worth US\$10, US\$15, US\$20, US\$25, and US\$35, respectively. Continued participation through the daily assessment was further incentivized by using a raffle for which an iPad mini was the prize. The completion of all seven surveys each week resulted in one entry into the raffle drawing.

Wikipedia mobile app

We assessed Wikipedia browsing data from 482,760 people collated across 50 countries or territories and 14 languages. Naturalistic Wikipedia browsing data was collected for the months of March 2022 and October 2022 from the webrequest logs⁵. In general, this data is continuously and automatically collected for analytic purposes on Wikimedia’s infrastructure and deleted after 90 days. We only considered requests from the mobile app to articles in the main namespace. We discard requests marked as automated traffic or coming from bots⁶. We identify unique readers via the `wmfuid`⁷ associated with the app installation (readers of the app have the option to turn off sending this ID with their requests). In order to protect readers’ privacy, we remove sensitive information in several steps: discarding all requests from countries with fewer than 500 unique readers or from countries included in the country protection list⁸); dropping IP, user agent, and fine-grained geospatial information; dropping timestamps of the requests by only keeping the time-order as well as the day and hour of all requests. Across both time points, we sampled 14,900 readers who were matched for browsing the same number of page views as in the laboratory data using propensity score matching with the package `psmpy`⁹.

⁵https://wikitech.wikimedia.org/wiki/Analytics/Data_Lake/Traffic/Webrequest

⁶https://wikitech.wikimedia.org/wiki/Analytics/Data_Lake/Traffic/BotDetection

⁷<https://wikitech.wikimedia.org/wiki/X-Analytics>

⁸https://wikitech.wikimedia.org/wiki/Analytics/Data_Lake/Edits/Geoeditors/Public#Country_Protection_List

⁹<https://pypi.org/project/psmpy/>

This procedure was motivated by the fact that network size, measured by the number of page views, is a known confound of most network metrics. Using larger sample sizes gains statistical power and sample variance at the cost of slightly poorer matching. Therefore, we also applied this matching procedure to better match page views in smaller datasets of 1,490 and 149 readers.

We applied the same matching procedure when constructing datasets for readers who accessed from different countries and for readers who used different language versions of Wikipedia. The datasets we constructed include the following:

- readers of English Wikipedia accessing from specific countries (10 countries with the largest number of unique readers): Canada, India, the United States of America, the Philippines, Australia, Italy, the United Kingdom, France, the Netherlands, and Germany.
- readers of 14 language versions of Wikipedia considered in a prior study covering different language families and taking into account the number and distribution of speakers worldwide: English, German, Chinese, Russian, Arabic, Japanese, Hungarian, Spanish, Bengali, Dutch, Romanian, Hebrew, Hindi, and Ukrainian [46].
- more granular datasets of readers who accessed a given Wikipedia language edition from a given country (74 language-country pairs from the 14 language versions above; for privacy reasons we include only pairs with more than 5,000 unique readers and where the country is not included in the country protection list). This process yields 50 different countries or territories: Algeria, Argentina, Australia, Austria, Bangladesh, Belgium, Brazil, Canada, Chile, Colombia, Czechia, Denmark, Finland, France, Germany, Ghana, Greece, Hong Kong, Hungary, India, Indonesia, Ireland, Israel, Italy, Japan, Kenya, Korea, Malaysia, Mexico, Morocco, Nepal, the Netherlands, New Zealand, Nigeria, Norway, Peru, the Philippines, Poland, Portugal, Romania, Serbia, South Africa, Spain, Sri Lanka, Sweden, Switzerland, Taiwan, Ukraine, the United Kingdom, and the United States of America.

Wikispeedia

Both the laboratory data and the naturalistic data on Wikipedia browsing characterize naturalistic information seeking. This type of naturalistic information seeking was unconstrained by experimenters. To compare and contrast naturalistic information with a more constrained type of targeted navigation, we used a previously collected dataset on an online task called “Wikispeedia” [49, 50]. The Wikispeedia task asks participants to begin from a given source article and navigate to a given target article using the shortest path of hyperlinks between the source and target articles. Articles were condensed into a miniature version of Wikipedia that includes 4,604 articles. The task was successfully completed by 14,246 participants and took an average of 2.64 ± 6.01 minutes. The data was downloaded from a [public repository](#).

Population Indicators of Education, Well-being, and Equality

We were interested in population-level education, well-being, and equality because these are outcomes that curiosity has been hypothesized to support. To measure population-level indices of these factors, we collated data from several sources on people from countries that were included in the Wikipedia naturalistic data.

Spatial navigation is a useful variable because the same mechanisms thought to underlie the spatial navigation important for foraging are thought to underlie the mental navigation important for information foraging [37, 66]. To index spatial navigation abilities, we evaluated previously collected data on non-verbal spatial navigation using a wayfinding task in 397,162 participants from 38 countries [113]. Participants played a mobile video game, “Sea Hero Quest,” in which they navigate a boat in search of sea creatures. Performance in this game is predictive of real-world navigation ability and clinical conditions in which spatial navigation is known to be impaired such as Alzheimer’s disease [114, 115]. Performance is also influenced by one’s environment [59], raising the question of how the digital environment relates to performance. The wayfinding task provided players with a map indicating a start location and the location of targets to find in a given order.

Next, we assess education, well-being, and inequality. Whereas we used daily diary methods to measure indicators of individual-level mood and well-being in the laboratory data, we operationalize dimensions of population-level mood and well-being using countrywide indicators from the Human Development Index (2021), Gender Inequality Index (2021), and World Happiness Report (2021-2022) [116, 117]. We used the Human Development Index and Gender Inequality Index. For each country, a “Health” dimension includes variables such as life expectancy at birth, maternal mortality and adolescent birth rates. A “Knowledge” dimension includes variables such as the expected years of schooling and the mean years of schooling. See the compiled [technical notes](#) for these indices for further definitions and data sources.

Positive and negative affect were originally measured by the Gallup World Report and compiled by the World Happiness Report. The general form for these affect questions was: “Did you experience the following feelings during a lot of the day yesterday?” The possible answers are yes (1) or no (0), resulting in country-level averages ranging from 0 to 1. Positive affect is defined as the average of three measures: laughter, enjoyment, and doing interesting things. Laughter was measured with the answer to the question “Did you smile or laugh a lot yesterday?”, enjoyment was measured with the answer to the question “Did you experience the following feelings during A LOT OF THE DAY yesterday? How about Enjoyment?”, and doing interesting things was measured with the answer to the question “Did you learn or do something interesting yesterday?”

Negative affect was defined as the average of three measures: worry, sadness, and anger. Worry was measured with the answer to the question “Did you experience the following feelings during A LOT OF THE DAY yesterday? How

about Worry?”, sadness with the answer to the question “Did you experience the following feelings during A LOT OF THE DAY yesterday? How about Sadness?”, and anger with the answer to the question “Did you experience the following feelings during A LOT OF THE DAY yesterday? How about Anger?”

Knowledge networks

We construct knowledge networks by treating each article as a node and hyperlinks between articles as edges. We defined undirected and binary edges between two nodes s and t with a value of 1 if there is a hyperlink from s to t or from t to s . We considered all hyperlinks among articles from the same corresponding month as the data. Specifically, for each dataset, we considered the hyperlink network of the respective Wikipedia language version at the end of the month of access. For this purpose we use the corresponding snapshot of the pagelinks-table¹⁰ and resolve redirects.

Analysis

Network metrics

We calculated 11 topological metrics for each network. These metrics include the network size, number of edges, density, clustering coefficient, degree, characteristic path length, global efficiency, coreness, minimum description length, modularity index, and number of groups (or modules). We calculate these 11 metrics for each dataset to result in an n (participants) by 11 (metrics) matrix characterizing network structure. This matrix is used to calculate metrics that characterize the tight networks emblematic of hunters and the loose networks emblematic of busybodies. Below, we describe each of the 11 metrics and how we aggregate selected metrics to characterize whether a person’s browsing is more like a hunter, a busybody, or a dancer. All metrics were calculated using the NetworkX, BCTPY, and graph-tool packages [118–120].

Size, edges, density, and degree. The network size was calculated by the number of nodes n . The number of edges m characterizes how connected articles were by hyperlinks, where busybodies tend to visit less connected articles than hunters. The network density is the fraction of existing edges out of all possible edges, where busybodies tend to construct looser networks with less density than hunters. The density d of the network is defined as:

$$d = \frac{2m}{n(n-1)}, \quad (1)$$

where m is the number of edges in the network and n is the number of nodes in the network. The degree k of a given node is the number of hyperlink connections of the given node to all other nodes, where busybodies tend to construct less connected networks than hunters.

¹⁰https://www.mediawiki.org/wiki/Manual:Pagelinks_table

Clustering coefficient. The clustering coefficient can be defined as the fraction of existing edges out of all possible edges in node triplets. The clustering of a node u is defined as:

$$c_u = \frac{2T(n)}{k(n)(k(n) - 1)}, \quad (2)$$

where $T(n)$ is the number of triangles through node n and $k(n)$ is the degree of n .

Characteristic path length. The average characteristic path length is the mean shortest path length between all pairs of nodes. The average characteristic path length is defined as:

$$a = \sum_{\substack{s,t \in V \\ s \neq t}} \frac{d(s,t)}{n(n-1)}, \quad (3)$$

where V is the set of nodes in the network, $d(s,t)$ is the shortest path from the source node s to the target node t , and n is the number of nodes.

Global efficiency. The global efficiency is the harmonic mean of the path distances, and can be defined as:

$$E = \frac{1}{n(n-1)} \sum_{i \neq j} \frac{1}{d_{ij}}, \quad (4)$$

where n is the number of nodes and d_{ij} is the shortest path distance between the nodes i and j .

Core-periphery. A network has a core-periphery structure if a set of core nodes and a set of periphery nodes exist such that core-to-core connections are most common, core-to-periphery connections are less common, and periphery-to-periphery connections are least common (if they exist at all). A core-periphery structure can be detected in a network by partitioning the network into a core group and a periphery group such that the number of core-group edges is maximized while the number of periphery-group edges is minimized [121]. The coreness can be defined as:

$$Q_C = \frac{1}{v_C} \left(\sum_{i,j \in C_c} (w_{ij} - \gamma_C \bar{w}) - \sum_{i,j \in C_p} (w_{ij} - \gamma_C \bar{w}) \right), \quad (5)$$

where C_c is the set of all nodes in the core, C_p is the set of all nodes in the periphery, w_{ij} is the weight between nodes i and j , \bar{w} is the average edge weight, γ_C is a parameter that adjusts the size of the core, and v_C is a normalization constant.

Modularity, partitions, and minimum description length. Modularity was estimated by fitting a hierarchical degree-corrected stochastic block model [122]. The fitted generative model contains a hierarchical grouping of nodes, wherein the number of levels in the hierarchy and the number of groups in each level are inferred automatically. We solve the inference problem that

maximizes the Bayesian posterior probability for the modules or partitions b such that:

$$P(\mathbf{b} \mid \mathbf{A}) = \frac{P(\mathbf{A} \mid \boldsymbol{\theta}, \mathbf{b})P(\boldsymbol{\theta}, \mathbf{b})}{P(\mathbf{A})}, \quad (6)$$

where A is the network's adjacency matrix and $\boldsymbol{\theta}$ are additional model parameters that control how the node partition affects the structure of the network.

Using the fitted model, we compute three metrics: the number of groups b , the modularity Q , and the minimum description length. The number of groups is given by the fitted model's partitioning of the nodes into b modules. Busybodies tend to have more modules than hunters.

The modularity Q is calculated as a proportion of edges within versus between groups on the lowest level of the hierarchy using Newman's (generalized) modularity metric:

$$Q = \frac{1}{2M} \sum_r m_{rr} - \frac{m_r^2}{2M}, \quad (7)$$

where m is the number of edges that fall between communities s and r , or twice the number of edges if $s = r$, and M is the total number of edges. Busybodies tend to construct looser networks with weaker within-module connections and hence less modularity than hunters.

The minimum description length [123] is the amount of information required to describe the data using our fitted stochastic block model. Maximizing the posterior probability of $P(\mathbf{b} \mid \mathbf{A})$ gives the minimum description length Σ defined as:

$$\Sigma = -\ln P(\mathbf{A} \mid \boldsymbol{\theta}, \mathbf{b}) - \ln P(\boldsymbol{\theta}, \mathbf{b}). \quad (8)$$

This metric captures the amount of information necessary to describe the data given the model as well as the amount of information to describe the model itself.

Busybody-hunter metric. We aggregated a few selected network measures above into a busybody-hunter metric. The busybody-hunter metric characterizes how loosely or tightly people construct knowledge networks as they browse Wikipedia. In prior work, the edge weight, clustering coefficient, global efficiency, and characteristic path length were related to the tight structures created by a hunter-like curiosity and the loose structures created by a busybody-like curiosity [41]. Hence, we defined an aggregate busybody/hunter score using the number of edges, clustering coefficient, global efficiency, and characteristic path length. Here we also include the number of edges instead of the edge weight because we are analyzing binary rather than weighted networks. The standardized values that characterize hunter-like exploration were summed, including the number of edges, clustering coefficient, and global efficiency. The standardized values that characterize busybody-like exploration were subtracted, including the characteristic path length. The busybody-hunter metric is then the sum of the hunter-like standardized values minus the

busybody-like standardized values. This aggregation is performed for the convenience of analysis and visualization, and throughout, we also consider the individual contributions of each metric.

Forward flow. Forward flow characterizes the momentous “leaps” in thought characteristic of creativity [62]. Forward flow is calculated by the average distance between the current “thought” and all previous thoughts. While forward flow of thoughts has been correlated with metrics of creativity, here we were interested in forward flow as a quality that we expect dancers to possess as they weave a thread through Wikipedia pages to seek information and build knowledge networks. Dancers are thought to take leaps of creative imagination, and hence we would expect dancer-like browsing to have higher forward flow than non-dancer-like browsing. Forward flow is calculated as:

$$\frac{\sum_{i=1}^{n-1} D_{n,i}}{n-1}, \quad (9)$$

where D is the distance between pages and n is the index of the page in the sequence of pages a reader has browsed. For this metric, we defined D using the cosine distance between two pages’ pre-trained word-vector embeddings from fastText [124].

Distance between distributions

To quantify the similarity between two samples for a given variable x (e.g., any of the network metrics above), we use the Kolmogorov-Smirnov (KS) distance. This distance is commonly used as a test statistic for the KS-test which assesses whether two empirical distributions are significantly different. Given two samples (of possibly different sizes n and m), $\{x_{1,i}\}_{i=1,\dots,n}$ and $\{x_{2,j}\}_{j=1,\dots,m}$, the KS-distance is defined as

$$\sup_x |F_{1,n}(x) - F_{2,m}(x)|, \quad (10)$$

where $F_{1,n}(x)$ and $F_{2,m}(x)$ are the empirical cumulative distribution functions and \sup is the supremum. The KS-distance yields values between 0 and 1.

Network models

For each knowledge network, we generated a set of corresponding random networks to compare the empirical data to simulated data with explicitly defined generative models. This comparison allows us to interpret the network structure of naturalistic browsing as a possible consequence of the rules that govern the growth and structure of networks according to generative models. Moreover, by comparing empirical to random networks we can determine how much our network measures are distinctive of empirical information seeking behavior versus random chance.

We used each network’s size, average node degree, and average edge density to parameterize generative models including randomly shuffled, Erdős-Rényi,

Barabási-Albert, random walks according to total popularity, random walks on local hyperlinks, a hybrid random walk based on a prior analysis of reader behavior [52], and a foraging model rooted in theories from ecology and psychology [76, 125].

Random networks. Randomly shuffled networks were constructed by randomizing the structure of the observed network by swapping edges between pairs of nodes while preserving the degree distribution [126]. The Erdős-Rényi network was generated using an edge probability corresponding to the edge density of the reader’s network size [127]. The Barabási-Albert model generates scale-free networks with heterogeneous degree distribution by a preferential attachment (rich-get-richer) mechanism [128].

We also created synthetic knowledge networks by simulating browsing as different types of random walks on the hyperlink network. For each synthetic network, we start random walks from the same first page and with the same number of steps as observed empirically. The next step of the random walks were determined by different rules. The random walks weave a thread through the underlying Wikipedia network according to four different generative rules.

Random walk. The first rule of determining the next step of the random walk was by sampling from a uniform probability distribution across the hyperlinks [129]. At each step, we pick one of the hyperlinks on the page at random, whereby each hyperlink has the same weight. If there were no hyperlinks, a page was picked at random proportional to the page’s popularity in terms of the number of page views accrued during the respective month. Random walks have been previously shown to resemble optimal foraging and thus serve as an important null model [51].

Random walk by popularity. The second rule of determining the next step of the random walk was by sampling from pages weighted by their popularity, forming a biased random walk similar to previous research [130]. Popularity was operationalized by the total number of page views during the month of access. Biased random walks based on popularity serve as a null model for the average behavior of Wikipedia readers.

Random walk by popularity and reader statistics. The third rule of determining the next step of the random walk was a mixture of the first and second rules. The next step would be determined by the first rule with a 37.5% probability or by the second rule with a 62.5% probability. This specific set of probabilities was drawn from prior work that found around 37.5% of page loads were from a different page within Wikipedia, while the remainder of the page loads were accessed from an external source [52].

Random walk by search and foraging. The fourth and final rule of determining the next step of the random walk was sampling pages according to a previously fit model of information search and foraging. In a prior study, this model was fit to the 149 participants of the laboratory data, producing an empirical distribution of 149 sets of two parameters: edge reinforcement and Lévy flight dynamics [39, 41]. Here we repeatedly draw pairs of parameters with replacement from the empirical distribution in order to simulate browsing. We

then test the similarity of structure between the simulated knowledge networks and the mobile app knowledge networks using the network metrics described above.

The edge reinforcement parameter represents a memory of familiarity that increases the probability of taking previously traversed paths. Mathematically, random walks across a path will increase the weighted value of that path, thereby increasing its future transition probability. When applied to human information seeking, edge reinforcement is associated with the personality trait of deprivation sensitivity [10, 42], a dimension of curiosity associated with aversion to uncertainty and gaps in knowledge.

The Lévy flight dynamics parameter characterizes an optimal distribution of step distances with many small steps and a few large leaps. Lévy flight dynamics are characterized by the manner in which the probability of a step decays as a function of distance [73]. The steepness of decay is directly related to the exponent defining the function's form. In particular, we measured the exponent of the decaying probability distribution by the step distance of the empirically observed behavior. Optimally efficient Lévy flight dynamics exist if this exponent is approximately 2 [125]. When applied to human information seeking, recent work shows that humans indeed exhibit an average exponent of 2.11 ± 0.15 suggestive of Lévy-like dynamics in curiosity-driven information seeking [41].

Information diversity

The diversity of information is important for conceptualizations of epistemic curiosity. This kind of curiosity is a drive for knowledge. In a historico-philosophical framework of curiosity styles, there are two complementary inclinations to seek knowledge. Hunters tend pursue a set of similar topics, whereas busybodies tend to explore widely in diverse topics. These tendencies had not been previously tested.

To calculate information diversity, we first sought to classify each article into distinct topics. We used a prior classification method that assigns each page to four topics: culture, geography, history and society, and science, technology, engineering and math (STEM) [57]. The broad topic "Culture" includes subtopics of visual arts, sports, philosophy and religion, performing arts, media, literature, linguistics, internet culture, food and drink, and biography. The broad topic "Geography" includes subtopics of specific regions and general geographical phenomena. The broad topic "History and Society" includes transportation, society, politics and government, military and welfare, history, education, and business/economics. Finally, the broad topic "STEM" includes technology, space, physics, medicine/health, mathematics, libraries/information, engineering, earth/environment, computing, chemistry, and biology.

Using this classification, we then calculated for each reader the proportion of page visits classified as culture, geography, history and society, or STEM.

The information diversity is a function of these proportions, based on the Shannon diversity index:

$$H = - \sum_i p_i \times \ln(p_i), \quad (11)$$

where p_i is the proportion of the pages made up of topic $i \in$ culture, geography, history and society, and STEM. The maximum information diversity occurs when people browse all topics evenly, whereas the minimum information diversity occurs when people browse only one topic.

Unsupervised clustering

We assess how knowledge networks cluster together using Uniform Manifold Approximation and Projection (UMAP) dimensionality reduction on the n -by-11 matrices of network structure [131]. We also include in this clustering several canonical random networks with explicitly known generative models. Clustering both empirical and random networks serves two purposes. First, their comparison allows us to determine how much our network measures can distinguish the structure of information seeking behavior from random chance. Here random networks serve as null models. Second, their comparison allows us to interpret how similar human behavior is to known generative principles, such as preferential attachment models (the rich-get-richer) and random walks. Finally, we used the 2-dimensional space generated by UMAP as a data-driven approach to detect different kinds of curiosity. To do so, we assessed the correlation between the coordinates of each dimension and either (i) a variable hypothesized to operationalize a busybody-hunter axis of curiosity or (ii) a variable hypothesized to operationalize a dancer axis of curiosity.

Citation Diversity Statement

Recent work in several fields of science has identified a bias in citation practices such that papers from women and other minority scholars are under-cited relative to the number of such papers in the field [132–140]. Here we sought to proactively consider choosing references that reflect the diversity of the field in thought, form of contribution, gender, race, ethnicity, and other factors. First, we obtained the predicted gender of the first and last author of each reference by using databases that store the probability of a first name being carried by a woman [136, 141]. By this measure (and excluding self-citations to the first and last authors of our current paper), our references contain 18.76% woman(first)/woman(last), 15.83% man/woman, 12.72% woman/man, and 52.7% 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. 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 [142, 143]. By this measure (and excluding self-citations), our references contain 12.23%

author of color (first)/author of color(last), 14.06% white author/author of color, 18.62% author of color/white author, and 55.10% 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.

Acknowledgments. We would like to thank Leila Zia for insightful discussions as well reviewing an initial draft of this paper. D.Z. acknowledges support from the George E. Hewitt Foundation for Medical Research. D.S.B. acknowledges support from the Center for Curiosity. D.M.L. acknowledges support from the National Institute on Drug Abuse (K01 DA047417).

Code and data availability. The code used in this analysis has been deposited to [a public repository](#). Server logs may contain sensitive information with implications for the privacy of readers. Aggregated data measures are included in the repository where privacy requirements are met.

Appendix A Supplementary Figures

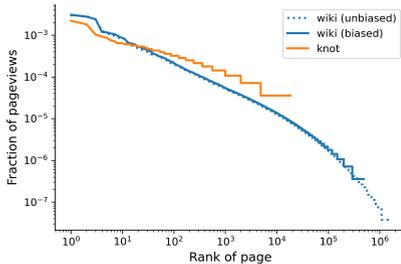


Fig. A1 Rank frequency plot of the number of times each distinct page was visited across all reading sessions for the different datasets.

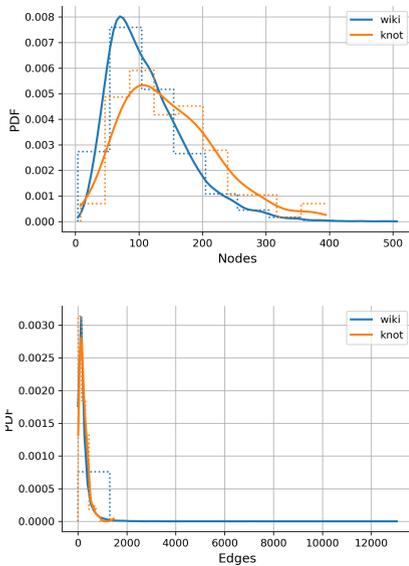


Fig. A2 Distribution of the network metrics in knowledge networks from Wikipedia and KNOT: number of nodes and number of edges. The solid line indicates the PDF from kernel density estimation; the dotted line indicates the normalized histograms.

Mean	Cluster 1	Cluster 2
N	2.116825e+06	5.178014e+05
E	1.455744e+08	4.939047e+07
k	6.679779e+01	8.482758e+01
ρ	1.771025e-04	4.479448e-04

Table A1 Differences between clusters driven by size. Cluster 1 includes German, English, Spanish, Hindi, Japanese, Dutch, Russian, Ukrainian, and Chinese language Wikipedias. Cluster 2 includes Arabic, Bengali, Hebrew, Hungarian, and Romanian language Wikipedias. The clusters split in dimensionality reduction due to cluster 2 having a smaller number of nodes, defined as articles. N: nodes; E: edges; k: degree; ρ : density.

References

- [1] Miller, G.A.: Informavores. In: Machlup, F., Mansfield, U. (eds.) *The Study of Information: Interdisciplinary Messages*, pp. 111–113. New York: Wiley, ??? (1983)
- [2] Pirolli, P., Card, S.: Information foraging. *Psychological review* **106**(4), 643 (1999)
- [3] Berlyne, D.E.: Curiosity and exploration: Animals spend much of their time seeking stimuli whose significance raises problems for psychology. *Science* **153**(3731), 25–33 (1966)
- [4] Berlyne, D.E., et al.: *A theory of human curiosity* (1954)
- [5] Silvia, P.J.: Interest—the curious emotion. *Current directions in psychological science* **17**(1), 57–60 (2008)
- [6] Kidd, C., Hayden, B.Y.: The psychology and neuroscience of curiosity. *Neuron* **88**(3), 449–460 (2015). <https://doi.org/10.1016/j.neuron.2015.09.010>
- [7] Kashdan, T.B., Disabato, D.J., Goodman, F.R., McKnight, P.E.: The Five-Dimensional curiosity scale revised (5DCR): Briefer subscales while separating overt and covert social curiosity. *Personality and individual differences* **157**, 109836 (2020). <https://doi.org/10.1016/j.paid.2020.109836>
- [8] Gottlieb, J., Oudeyer, P.-Y., Lopes, M., Baranes, A.: Information-seeking, curiosity, and attention: computational and neural mechanisms. *Trends in cognitive sciences* **17**(11), 585–593 (2013). <https://doi.org/10.1016/j.tics.2013.09.001>
- [9] Gottlieb, J., Oudeyer, P.-Y.: Towards a neuroscience of active sampling and curiosity. *Nature Reviews Neuroscience* **19**(12), 758–770 (2018)
- [10] Kashdan, T.B., Stikma, M.C., Disabato, D.J., McKnight, P.E., Bekier, J., Kaji, J., Lazarus, R.: The five-dimensional curiosity scale: Capturing the bandwidth of curiosity and identifying four unique subgroups of curious people. *Journal of Research in Personality* **73**, 130–149 (2018)
- [11] Zurn, P.: Busybody, hunter, dancer: three historical models of curiosity. Toward new philosophical explorations of the epistemic desire to know: Just curious about curiosity, 27–49 (2019)
- [12] Zurn, P., Bassett, D.S.: *Curious Minds: The Power of Connection*. MIT Press, ??? (2022)

- [13] Blasi, D.E., Henrich, J., Adamou, E., Kemmerer, D., Majid, A.: Over-reliance on english hinders cognitive science. *Trends in cognitive sciences* **26**(12), 1153–1170 (2022)
- [14] Henrich, J., Heine, S.J., Norenzayan, A.: The weirdest people in the world? *Behavioral and brain sciences* **33**(2-3), 61–83 (2010)
- [15] Zurn, P., Bassett, D.S.: On curiosity: A fundamental aspect of personality, a practice of network growth. *Personality neuroscience* **1**, 13 (2018). <https://doi.org/10.1017/pen.2018.3>
- [16] Medina, J.: *The Epistemology of Resistance: Gender and Racial Oppression, Epistemic Injustice, and Resistant Imaginations*. Oxford University Press, ??? (2012)
- [17] Zurn, P.: *Curiosity and Power: The Politics of Inquiry*. U of Minnesota Press, ??? (2021)
- [18] Zurn, P.: Curiosity: An affect of resistance. *Theory & Event* **24**(2), 611–617 (2021)
- [19] Park, N., Peterson, C., Seligman, M.E.: Strengths of character and well-being. *Journal of social and Clinical Psychology* **23**(5), 603–619 (2004)
- [20] Kang, M.J., Hsu, M., Krajbich, I.M., Loewenstein, G., McClure, S.M., Wang, J.T.-y., Camerer, C.F.: The wick in the candle of learning: Epistemic curiosity activates reward circuitry and enhances memory. *Psychological science* **20**(8), 963–973 (2009)
- [21] Gruber, M.J., Gelman, B.D., Ranganath, C.: States of curiosity modulate hippocampus-dependent learning via the dopaminergic circuit. *Neuron* **84**(2), 486–496 (2014)
- [22] Wade, S., Kidd, C.: The role of prior knowledge and curiosity in learning. *Psychonomic bulletin & review* **26**(4), 1377–1387 (2019). <https://doi.org/10.3758/s13423-019-01598-6>
- [23] Freire, P.: *Pedagogy of the oppressed*. In: *Toward a Sociology of Education*, pp. 374–386. Routledge, ??? (2020)
- [24] Clark, J., Vincent, A., Wang, X., McGowan, A.L., Lydon-Staley, D.M.: Smokers’ curiosity facilitates recall of tobacco-related health information. *Health Communication*, 1–9 (2022)
- [25] Lyew, T., Ikhlas, A., Sayed, F., Vincent, A., Lydon-Staley, D.: Curiosity, surprise, and the recall of tobacco-related health information in adolescents. *Journal of Health Communication*, 1–12 (2023)

- [26] Zedelius, C.M., Gross, M.E., Schooler, J.W.: Inquisitive but not discerning: Deprivation curiosity is associated with excessive openness to inaccurate information. *Journal of Research in Personality* **98**, 104227 (2022)
- [27] Schiefele, U.: Interest and learning from text. *Scientific studies of reading* **3**(3), 257–279 (1999)
- [28] Renner, B.: Curiosity about people: The development of a social curiosity measure in adults. *Journal of personality assessment* **87**(3), 305–316 (2006)
- [29] Kaczmarek, L.D., Baczkowski, B., Enko, J., Baran, B., Theuns, P.: Subjective well-being as a mediator for curiosity and depression. *Polish Psychological Bulletin* **45**(2), 200–204 (2014)
- [30] Lydon-Staley, D.M., Zurn, P., Bassett, D.S.: Within-person variability in curiosity during daily life and associations with well-being. *Journal of personality* **88**(4), 625–641 (2020)
- [31] Silvia, P.J., Christensen, A.P.: Looking up at the curious personality: Individual differences in curiosity and openness to experience. *Current Opinion in Behavioral Sciences* **35**, 1–6 (2020)
- [32] Kristeva, J.: *Black Sun: Depression and Melancholia*. Columbia University Press, ??? (1989)
- [33] Steglich-Petersen, A., Varga, S.: Curiosity and zetetic style in adhd. *Philosophical Psychology*, 1–25 (2023)
- [34] Lindgren, K.P., Mullins, P.M., Neighbors, C., Blayney, J.A.: Curiosity killed the cocktail? curiosity, sensation seeking, and alcohol-related problems in college women. *Addictive Behaviors* **35**(5), 513–516 (2010)
- [35] Loewenstein, G.: The psychology of curiosity: A review and reinterpretation. *Psychological bulletin* **116**(1), 75–98 (1994). <https://doi.org/10.1037/0033-2909.116.1.75>
- [36] Li, T., Huang, H., Liu, J., Tang, X.: Killing the cats or satisfying the human? the role of epistemic curiosity in adolescents' multidimensional well-being. *Journal of Pacific Rim Psychology* **17**, 18344909231185381 (2023)
- [37] Todd, P.M., Hills, T.T.: Foraging in mind. *Current Directions in Psychological Science* **29**(3), 309–315 (2020)
- [38] Bassett, D.S.: A network science of the practice of curiosity. In: Zurn,

- P., Shankar, A. (eds.) *Curiosity Studies: A New Ecology of Knowledge*. University of Minnesota Press, ??? (2020)
- [39] Zhou, D., Lydon-Staley, D.M., Zurn, P., Bassett, D.S.: The growth and form of knowledge networks by kinesthetic curiosity. *Current opinion in behavioral sciences* **35**, 125–134 (2020)
- [40] Patankar, S.P., Zhou, D., Lynn, C.W., Kim, J.Z., Ouellet, M., Ju, H., Zurn, P., Lydon-Staley, D.M., Bassett, D.S.: Curiosity as filling, compressing, and reconfiguring knowledge networks. arXiv preprint arXiv:2204.01182 (2022)
- [41] Lydon-Staley, D.M., Zhou, D., Blevins, A.S., Zurn, P., Bassett, D.S.: Hunters, busybodies and the knowledge network building associated with deprivation curiosity. *Nature human behaviour* **5**(3), 327–336 (2021)
- [42] Litman, J.A., Jimerson, T.L.: The measurement of curiosity as a feeling of deprivation. *Journal of personality assessment* **82**(2), 147–157 (2004)
- [43] Wikistats - statistics for Wikimedia Projects. <https://stats.wikimedia.org/>. Accessed: 2023-7-19
- [44] Domantas, G.: The most visited website in every country (that isn't A search engine). Hostinger. Accessed: 2023-7-18 (2022)
- [45] Singer, P., Lemmerich, F., West, R., Zia, L., Wulczyn, E., Strohmaier, M., Leskovec, J.: Why we read wikipedia. In: *Proceedings of the 26th International Conference on World Wide Web*, pp. 1591–1600 (2017)
- [46] Lemmerich, F., Sáez-Trumper, D., West, R., Zia, L.: Why the world reads wikipedia: Beyond english speakers. In: *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*, pp. 618–626 (2019)
- [47] Hill, B.M., Shaw, A.: The most important laboratory for social scientific and computing research in history. In: *Wikipedia @ 20*, pp. 159–174. The MIT Press, ??? (2020). <https://doi.org/10.7551/mitpress/12366.003.0015>
- [48] Rosenbaum, P.R., Rubin, D.B.: The central role of the propensity score in observational studies for causal effects. *Biometrika* **70**(1), 41–55 (1983). <https://doi.org/10.1093/biomet/70.1.41>
- [49] West, R., Pineau, J., Precup, D.: Wikispeedia: An online game for inferring semantic distances between concepts. In: *Twenty-First International Joint Conference on Artificial Intelligence* (2009)

- [50] West, R., Leskovec, J.: Human wayfinding in information networks. In: Proceedings of the 21st International Conference on World Wide Web, pp. 619–628 (2012)
- [51] Abbott, J.T., Austerweil, J.L., Griffiths, T.L.: Random walks on semantic networks can resemble optimal foraging. In: Neural Information Processing Systems Conference; A Preliminary Version of this Work Was Presented at the Aforementioned Conference., vol. 122, p. 558 (2015). American Psychological Association
- [52] Piccardi, T., Gerlach, M., Arora, A., West, R.: A large-scale characterization of how readers browse wikipedia. *ACM Transactions on the Web* **17**(2), 1–22 (2023)
- [53] Keller, A.S., Leikauf, J.E., Holt-Gosselin, B., Staveland, B.R., Williams, L.M.: Paying attention to attention in depression. *Translational psychiatry* **9**(1), 279 (2019)
- [54] Whitmer, A.J., Gotlib, I.H.: An attentional scope model of rumination. *Psychological bulletin* **139**(5), 1036 (2013)
- [55] Fredrickson, B.L.: The role of positive emotions in positive psychology: The broaden-and-build theory of positive emotions. *American psychologist* **56**(3), 218 (2001)
- [56] Kelly, C.A., Sharot, T.: Individual differences in information-seeking. *Nature communications* **12**(1), 7062 (2021)
- [57] Johnson, I., Gerlach, M., Sáez-Trumper, D.: Language-agnostic topic classification for wikipedia. In: Companion Proceedings of the Web Conference 2021, pp. 594–601 (2021)
- [58] Van Praag, H., Kempermann, G., Gage, F.H.: Neural consequences of environmental enrichment. *Nature Reviews Neuroscience* **1**(3), 191–198 (2000)
- [59] Coutrot, A., Silva, R., Manley, E., de Cothi, W., Sami, S., Bohbot, V.D., Wiener, J.M., Hölscher, C., Dalton, R.C., Hornberger, M., *et al.*: Global determinants of navigation ability. *Current Biology* **28**(17), 2861–2866 (2018)
- [60] Lydon-Staley, D.M., Falk, E.B., Bassett, D.S.: Within-person variability in sensation-seeking during daily life: Positive associations with alcohol use and self-defined risky behaviors. *Psychology of addictive behaviors* **34**(2), 257 (2020)
- [61] Tversky, B., Hemenway, K.: Objects, parts, and categories. *Journal of*

- experimental psychology: *General* **113**(2), 169 (1984)
- [62] Gray, K., Anderson, S., Chen, E.E., Kelly, J.M., Christian, M.S., Patrick, J., Huang, L., Kenett, Y.N., Lewis, K.: “forward flow”: A new measure to quantify free thought and predict creativity. *American Psychologist* **74**(5), 539 (2019)
- [63] Kenett, Y.N., Anaki, D., Faust, M.: Investigating the structure of semantic networks in low and high creative persons. *Frontiers in human neuroscience* **8**, 407 (2014)
- [64] Kenett, Y.N., Faust, M.: A semantic network cartography of the creative mind. *Trends in cognitive sciences* **23**(4), 271–274 (2019)
- [65] Replications do not fail. *Nature human behaviour* **4**(6), 559 (2020). <https://doi.org/10.1038/s41562-020-0903-0>
- [66] Bellmund, J.L., Gärdenfors, P., Moser, E.I., Doeller, C.F.: Navigating cognition: Spatial codes for human thinking. *Science* **362**(6415), 6766 (2018)
- [67] Aru, J., Drüke, M., Pikamäe, J., Larkum, M.E.: Mental navigation and the neural mechanisms of insight. *Trends in Neurosciences* **46**(2), 100–109 (2023)
- [68] Schiller, D., Eichenbaum, H., Buffalo, E.A., Davachi, L., Foster, D.J., Leutgeb, S., Ranganath, C.: Memory and space: towards an understanding of the cognitive map. *Journal of Neuroscience* **35**(41), 13904–13911 (2015)
- [69] Epstein, R.A., Patai, E.Z., Julian, J.B., Spiers, H.J.: The cognitive map in humans: spatial navigation and beyond. *Nature neuroscience* **20**(11), 1504–1513 (2017)
- [70] Momennejad, I.: Learning structures: predictive representations, replay, and generalization. *Current Opinion in Behavioral Sciences* **32**, 155–166 (2020)
- [71] Morton, N.W., Preston, A.R.: Concept formation as a computational cognitive process. *Current Opinion in Behavioral Sciences* **38**, 83–89 (2021)
- [72] Christoff, K., Irving, Z.C., Fox, K.C., Spreng, R.N., Andrews-Hanna, J.R.: Mind-wandering as spontaneous thought: a dynamic framework. *Nature reviews neuroscience* **17**(11), 718–731 (2016)
- [73] Viswanathan, G.M., Da Luz, M.G., Raposo, E.P., Stanley, H.E.: The

- Physics of Foraging: an Introduction to Random Searches and Biological Encounters. Cambridge University Press, ??? (2011)
- [74] Rhodes, T., Turvey, M.T.: Human memory retrieval as lévy foraging. *Physica A: Statistical Mechanics and its Applications* **385**(1), 255–260 (2007)
- [75] Zhu, J., Sanborn, A., Chater, N.: Mental sampling in multimodal representations. *Advances in neural information processing systems* **31** (2018)
- [76] Garg, K., Kello, C.T.: Efficient lévy walks in virtual human foraging. *Scientific reports* **11**(1), 5242 (2021)
- [77] Fredrickson, B.L.: Positive emotions broaden and build. In: *Advances in Experimental Social Psychology* vol. 47, pp. 1–53. Elsevier, ??? (2013)
- [78] Hsiung, A., Poh, J.-H., Huettel, S.A., Adcock, R.A.: Curiosity evolves as information unfolds. *Proceedings of the National Academy of Sciences* **120**(43), 2301974120 (2023) <https://www.pnas.org/doi/pdf/10.1073/pnas.2301974120>. <https://doi.org/10.1073/pnas.2301974120>
- [79] Kruglanski, A.W., Webster, D.M.: Motivated closing of the mind: “seizing” and “freezing”. *The motivated mind*, 60–103 (2018)
- [80] Lynn, C.W., Bassett, D.S.: Quantifying the compressibility of complex networks. *Proceedings of the National Academy of Sciences* **118**(32), 2023473118 (2021)
- [81] Boguna, M., Krioukov, D., Claffy, K.C.: Navigability of complex networks. *Nature Physics* **5**(1), 74–80 (2009)
- [82] Garg, K., Padilla-Iglesias, C., Restrepo Ochoa, N., Knight, V.B.: Hunter–gatherer foraging networks promote information transmission. *Royal Society Open Science* **8**(12), 211324 (2021)
- [83] Bartumeus, F., Peters, F., Pueyo, S., Marrasé, C., Catalan, J.: Helical lévy walks: adjusting searching statistics to resource availability in microzooplankton. *Proceedings of the National Academy of Sciences* **100**(22), 12771–12775 (2003)
- [84] Dubey, R., Griffiths, T.L.: Reconciling novelty and complexity through a rational analysis of curiosity. *Psychological Review* **127**(3), 455 (2020)
- [85] Gopnik, A.: Childhood as a solution to explore–exploit tensions. *Philosophical Transactions of the Royal Society B* **375**(1803), 20190502

(2020)

- [86] Harhen, N.C., Bornstein, A.M.: Overharvesting in human patch foraging reflects rational structure learning and adaptive planning. *Proceedings of the National Academy of Sciences* **120**(13), 2216524120 (2023)
- [87] Hong, L., Lalmas, M.: Tutorial on online user engagement: Metrics and optimization. In: *Companion Proceedings of The 2019 World Wide Web Conference. WWW '19*, pp. 1303–1305. Association for Computing Machinery, New York, NY, USA (2019). <https://doi.org/10.1145/3308560.3320087>
- [88] TeBlunthuis, N., Bayer, T., Vasileva, O.: Dwelling on wikipedia: investigating time spent by global encyclopedia readers. In: *Proceedings of the 15th International Symposium on Open Collaboration*, p. 14. ACM, ??? (2019). <https://doi.org/10.1145/3306446.3340829>
- [89] Sáez-Trumper, D.: Disinformation and AI: The differences between Wikipedia and social media. <https://diff.wikimedia.org/2021/09/15/disinformation-and-ai-the-differences-between-wikipedia-and-social-media/>. Accessed: 2022-12-8 (2021)
- [90] Wineburg, S., McGrew, S.: Lateral reading and the nature of expertise: Reading less and learning more when evaluating digital information. *Teachers College Record* **121**(11), 1–40 (2019). <https://doi.org/10.1177/016146811912101102>
- [91] Patankar, S.P., Ouellet, M., Cervino, J., Ribeiro, A., Murphy, K.A., Bassett, D.S.: Intrinsically motivated graph exploration using network theories of human curiosity (2023) [arXiv:2307.04962](https://arxiv.org/abs/2307.04962) [cs.LG]
- [92] Lee, H., Chen, J.: Predicting memory from the network structure of naturalistic events. *Nature Communications* **13**(1), 4235 (2022)
- [93] Lamprecht, D., Lerman, K., Helic, D., Strohmaier, M.: How the structure of wikipedia articles influences user navigation. *New Review of Hypermedia and Multimedia* **23**(1), 29–50 (2017)
- [94] Rodi, G.C., Loreto, V., Tria, F.: Search strategies of wikipedia readers. *PloS one* **12**(2), 0170746 (2017)
- [95] Ruggeri, K., Većkalov, B., Bojanić, L., Andersen, T.L., Ashcroft-Jones, S., Ayacaxli, N., Barea-Arroyo, P., Berge, M.L., Bjørndal, L.D., Bursalioglu, A., *et al.*: The general fault in our fault lines. *Nature Human Behaviour* **5**(10), 1369–1380 (2021)
- [96] Kelly, C., Sharot, T.: Knowledge-seeking reflects and shapes well-being

(2023)

- [97] Brady, W.J., Jackson, J.C., Lindström, B., Crockett, M.: Algorithm-mediated social learning in online social networks. Preprint at OSF preprints. <https://doi.org/10.31219/osf.io/yw5ah> (2023)
- [98] Ram, N., Yang, X., Cho, M.-J., Brinberg, M., Muirhead, F., Reeves, B., Robinson, T.N.: Screenomics: A new approach for observing and studying individuals' digital lives. *Journal of Adolescent Research* **35**(1), 16–50 (2020)
- [99] Brinberg, M., Ram, N., Yang, X., Cho, M.-J., Sundar, S.S., Robinson, T.N., Reeves, B.: The idiosyncrasies of everyday digital lives: Using the human screenome project to study user behavior on smartphones. *Computers in Human Behavior* **114**, 106570 (2021)
- [100] Collins, P.H.: *Intersectionality as Critical Social Theory*. Duke University Press, ??? (2019)
- [101] Kidd, I., James, J.M., Jr, G.P. (eds.): *The Routledge Handbook of Epistemic Injustice*. Taylor and Francis, ??? (2017)
- [102] Manne, K.: *Down Girl: The Logic of Misogyny*. Oxford University Press, ??? (2017)
- [103] Perry, I.: *Vexy Thing*. Duke University Press, ??? (2018)
- [104] Collins, P.H.: *The routledge handbook of epistemic injustice*, pp. 115–124. Routledge (2017)
- [105] Box-Steffensmeier, J.M., Burgess, J., Corbetta, M., Crawford, K., Duflo, E., Fogarty, L., Gopnik, A., Hanafi, S., Herrero, M., Hong, Y.-Y., Kameyama, Y., Lee, T.M.C., Leung, G.M., Nagin, D.S., Nobre, A.C., Nordentoft, M., Okbay, A., Perfors, A., Rival, L.M., Sugimoto, C.R., Tungodden, B., Wagner, C.: The future of human behaviour research. *Nature human behaviour* **6**(1), 15–24 (2022). <https://doi.org/10.1038/s41562-021-01275-6>
- [106] Edelman, A., Wolff, T., Montagne, D., Bail, C.A.: Computational social science and sociology. *Annual review of sociology* **46**(1), 61–81 (2020). <https://doi.org/10.1146/annurev-soc-121919-054621>
- [107] Lazer, D., Hargittai, E., Freelon, D., Gonzalez-Bailon, S., Munger, K., Ognyanova, K., Radford, J.: Meaningful measures of human society in the twenty-first century. *Nature* **595**(7866), 189–196 (2021). <https://doi.org/10.1038/s41586-021-03660-7>

- [108] Lazer, D.M.J., Pentland, A., Watts, D.J., Aral, S., Athey, S., Contractor, N., Freelon, D., Gonzalez-Bailon, S., King, G., Margetts, H., Nelson, A., Salganik, M.J., Strohmaier, M., Vespignani, A., Wagner, C.: Computational social science: Obstacles and opportunities. *Science* **369**(6507), 1060–1062 (2020). <https://doi.org/10.1126/science.aaz8170>
- [109] Jach, H., Cools, R., Frisvold, A., Grubb, M., Hartley, C., Hartmann, J., Hunter, L., Jia, R., de Lange, F., Larisch, R., et al.: Curiosity in cognitive science and personality psychology: Individual differences in information demand have a low dimensional structure that is predicted by personality traits (2023)
- [110] Sagar, A.D., Najam, A.: The human development index: a critical review. *Ecological economics* **25**(3), 249–264 (1998)
- [111] Alkire, S.: Dimensions of human development. *World development* **30**(2), 181–205 (2002)
- [112] Terry, P.C., Lane, A.M., Fogarty, G.J.: Construct validity of the profile of mood states—adolescents for use with adults. *Psychology of sport and exercise* **4**(2), 125–139 (2003)
- [113] Coutrot, A., Manley, E., Goodroe, S., Gahnstrom, C., Filomena, G., Yesiltepe, D., Dalton, R.C., Wiener, J.M., Hölscher, C., Hornberger, M., et al.: Entropy of city street networks linked to future spatial navigation ability. *Nature* **604**(7904), 104–110 (2022)
- [114] Coutrot, A., Schmidt, S., Coutrot, L., Pittman, J., Hong, L., Wiener, J.M., Hölscher, C., Dalton, R.C., Hornberger, M., Spiers, H.J.: Virtual navigation tested on a mobile app is predictive of real-world wayfinding navigation performance. *PloS one* **14**(3), 0213272 (2019)
- [115] Coughlan, G., Coutrot, A., Khondoker, M., Minihane, A.-M., Spiers, H., Hornberger, M.: Toward personalized cognitive diagnostics of at-genetic-risk alzheimer’s disease. *Proceedings of the National Academy of Sciences* **116**(19), 9285–9292 (2019)
- [116] Rowan, A.N.: World happiness report 2023. *WellBeing News* **5**(3), 1 (2023)
- [117] Programme), U.U.N.D.: Human development report 2021-22. UNDP (United Nations Development Programme) (2022)
- [118] Hagberg, A., Swart, P., S Chult, D.: Exploring network structure, dynamics, and function using networkx. Technical report, Los Alamos National Lab.(LANL), Los Alamos, NM (United States) (2008)

- [119] Rubinov, M., Sporns, O.: Complex network measures of brain connectivity: uses and interpretations. *Neuroimage* **52**(3), 1059–1069 (2010)
- [120] Peixoto, T.P.: The graph-tool python library. figshare (2014). <https://doi.org/10.6084/m9.figshare.1164194>. Accessed 2014-09-10
- [121] Borgatti, S.P., Everett, M.G.: Models of core/periphery structures. *Social networks* **21**(4), 375–395 (2000)
- [122] Peixoto, T.P.: Hierarchical block structures and High-Resolution model selection in large networks. *Physical Review X* **4**(1), 011047 (2014). <https://doi.org/10.1103/PhysRevX.4.011047>
- [123] Grünwald, P.D.: *The Minimum Description Length Principle*. MIT Press, Cambridge, MA (2007)
- [124] Joulin, A., Grave, E., Bojanowski, P., Douze, M., Jégou, H., Mikolov, T.: Fasttext.zip: Compressing text classification models. arXiv preprint arXiv:1612.03651 (2016)
- [125] Viswanathan, G., Afanasyev, V., Buldyrev, S.V., Havlin, S., Da Luz, M., Raposo, E., Stanley, H.E.: Lévy flights in random searches. *Physica A: Statistical Mechanics and its Applications* **282**(1-2), 1–12 (2000)
- [126] Maslov, S., Sneppen, K.: Specificity and stability in topology of protein networks. *Science* **296**(5569), 910–913 (2002)
- [127] ERDdS, P., R&wi, A.: On random graphs i. *Publ. math. debrecen* **6**(290-297), 18 (1959)
- [128] Albert, R., Barabási, A.-L.: Statistical mechanics of complex networks. *Reviews of modern physics* **74**(1), 47 (2002)
- [129] Révész, P.: *Random Walk in Random and Non-random Environments*. World Scientific, ??? (2013)
- [130] Sinatra, R., Gómez-Gardenes, J., Lambiotte, R., Nicosia, V., Latora, V.: Maximal-entropy random walks in complex networks with limited information. *Physical Review E* **83**(3), 030103 (2011)
- [131] McInnes, L., Healy, J., Melville, J.: Umap: Uniform manifold approximation and projection for dimension reduction. arXiv preprint arXiv:1802.03426 (2018)
- [132] Mitchell, S.M., Lange, S., Brus, H.: Gendered citation patterns in international relations journals. *International Studies Perspectives* **14**(4), 485–492 (2013)

- [133] Dion, M.L., Sumner, J.L., Mitchell, S.M.: Gendered citation patterns across political science and social science methodology fields. *Political Analysis* **26**(3), 312–327 (2018)
- [134] Caplar, N., Tacchella, S., Birrer, S.: Quantitative evaluation of gender bias in astronomical publications from citation counts. *Nature Astronomy* **1**(6), 0141 (2017)
- [135] Maliniak, D., Powers, R., Walter, B.F.: The gender citation gap in international relations. *International Organization* **67**(4), 889–922 (2013)
- [136] Dworkin, J.D., Linn, K.A., Teich, E.G., Zurn, P., Shinohara, R.T., Bassett, D.S.: The extent and drivers of gender imbalance in neuroscience reference lists. *bioRxiv* (2020) <https://www.biorxiv.org/content/early/2020/01/11/2020.01.03.894378.full.pdf>. <https://doi.org/10.1101/2020.01.03.894378>
- [137] Bertolero, M.A., Dworkin, J.D., David, S.U., Lloreda, C.L., Srivastava, P., Stiso, J., Zhou, D., Dzirasa, K., Fair, D.A., Kaczkurkin, A.N., Marlin, B.J., Shohamy, D., Uddin, L.Q., Zurn, P., Bassett, D.S.: Racial and ethnic imbalance in neuroscience reference lists and intersections with gender. *bioRxiv* (2020)
- [138] Wang, X., Dworkin, J.D., Zhou, D., Stiso, J., Falk, E.B., Bassett, D.S., Zurn, P., Lydon-Staley, D.M.: Gendered citation practices in the field of communication. *Annals of the International Communication Association* (2021). <https://doi.org/10.1080/23808985.2021.1960180>
- [139] Chatterjee, P., Werner, R.M.: Gender disparity in citations in high-impact journal articles. *JAMA Netw Open* **4**(7), 2114509 (2021)
- [140] Fulvio, J.M., Akinnola, I., Postle, B.R.: Gender (im)balance in citation practices in cognitive neuroscience. *J Cogn Neurosci* **33**(1), 3–7 (2021)
- [141] Zhou, D., Cornblath, E.J., Stiso, J., Teich, E.G., Dworkin, J.D., Blevins, A.S., Bassett, D.S.: Gender Diversity Statement and Code Notebook V1.0. <https://doi.org/10.5281/zenodo.3672110>. <https://doi.org/10.5281/zenodo.3672110>
- [142] Ambekar, A., Ward, C., Mohammed, J., Male, S., Skiena, S.: 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)
- [143] Sood, G., Laohaprapanon, S.: Predicting race and ethnicity from the sequence of characters in a name. *arXiv preprint arXiv:1805.02109* (2018)