

NETWORK SCIENCE

Architectural styles of curiosity in global Wikipedia mobile app readership

Dale Zhou^{1,2}, Shubhankar Patankar¹, David M. Lydon-Staley^{1,3}, Perry Zurn⁴, Martin Gerlach^{5*}†, Dani S. Bassett^{1,2,6,7,8,9,10*}†

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 analyze a naturalistic population of 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 replicate two styles of curiosity previously identified in laboratory studies: the nomadic "busybody" and the targeted "hunter." Further, we find evidence for another style—the "dancer"—which was previously predicted by a historico-philosophical examination of texts over two millennia and is characterized by creative modes of knowledge production. We identify associations, globally, between the structure of knowledge networks and population-level indicators of spatial navigation, education, mood, well-being, and inequality. These results advance our understanding of Wikipedia's global readership and demonstrate how cultural and geographical properties of the digital environment relate to different styles of curiosity.

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 multifaceted 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 overreliance on studies using the English language (13), as well as WEIRD (Western, Educated, Industrialized, Rich, and Democratic) populations (14). 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 (20–26). There is also evidence that, within the online ecosystem, curious individuals are better able to critically assess the novelty and quality of false information (27), to read deeply (28), and to explore others' perspectives, thoughts, and feelings (29). More generally, curiosity has been associated with individual well-being, lower depressed mood, and greater openness to experiences (10, 30–32). However, specific styles of curiosity may also be associated with depression, anxiety, and attention-deficit/hyperactivity disorder (33, 34). Moreover, curiosity can be double-edged, with high curiosity associated with behaving erroneously yet overconfidently, being less discerning of information quality, lacking intellectual humility (27), and having more addictive behavior (35). 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, which is consistently positively associated with indicators of well-being (10, 36). 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 (37)—and well-being (38). Thus, investigating the advantages and disadvantages

¹Department of Bioengineering, University of Pennsylvania, 240 Skirkanich Hall, Philadelphia, PA 19104, USA. ²Neuroscience Graduate Group, Perelman School of Medicine, University of Pennsylvania, 421 Curie Boulevard, Philadelphia, PA 19104, USA. ³Annenberg School of Communication, University of Pennsylvania, 3620 Walnut St, Philadelphia, PA 19104, USA. ⁴Department of Philosophy, American University, 4400 Massachusetts Ave NW, Washington, DC 20016, USA. ⁵Wikimedia Foundation, 1 Montgomery St, San Francisco, CA 94104, USA. ⁶Department of Electrical and Systems Engineering, University of Pennsylvania, 200 S 33rd St, Philadelphia, PA 19104, USA. ⁷Department of Neurology, University of Pennsylvania, 3400 Spruce St, Philadelphia, PA 19104, USA. ⁸Department of Psychiatry, University of Pennsylvania, 800 Spruce St, Philadelphia, PA 19104, USA. ⁹Santa Fe Institute, 1399 Hyde Park Rd, Santa Fe, NM 87501, USA. ¹⁰Montreal Neurological Institute, McGill University, Montreal, QC H3A 2B4, Canada.

*Corresponding author. Email: mgerlach@wikimedia.org (M.G.); dsb@seas.upenn.edu. (D.S.B.)

†These authors contributed equally to this work.

of different styles of curious practice is key to understanding human behavior writ large (26, 34).

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, 39). In this framework, the behavioral expression of curiosity is characterized by the construction of knowledge networks by individual readers (40). 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 (41, 42) 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 min a day for 21 days (36). 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, 37, 43).

Here, we generalize and expand the analysis of knowledge network structure by building a large, naturalistic dataset consisting of 482,760 Wikipedia readers accessing 14 different language editions from 50 countries or territories (Fig. 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 more than 60 million free articles in more than 300 languages (44). Wikipedia is the most popular website (alongside search engines, Facebook, and YouTube) in 43 countries, outpacing any other website (45), attracting 1.5 billion unique devices and billions of pageloads every month (44). People turn to Wikipedia when they need to learn about novel, uncertain, and complex information (46, 47). 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 that is one of the most important laboratories for social scientific research in history (48).

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

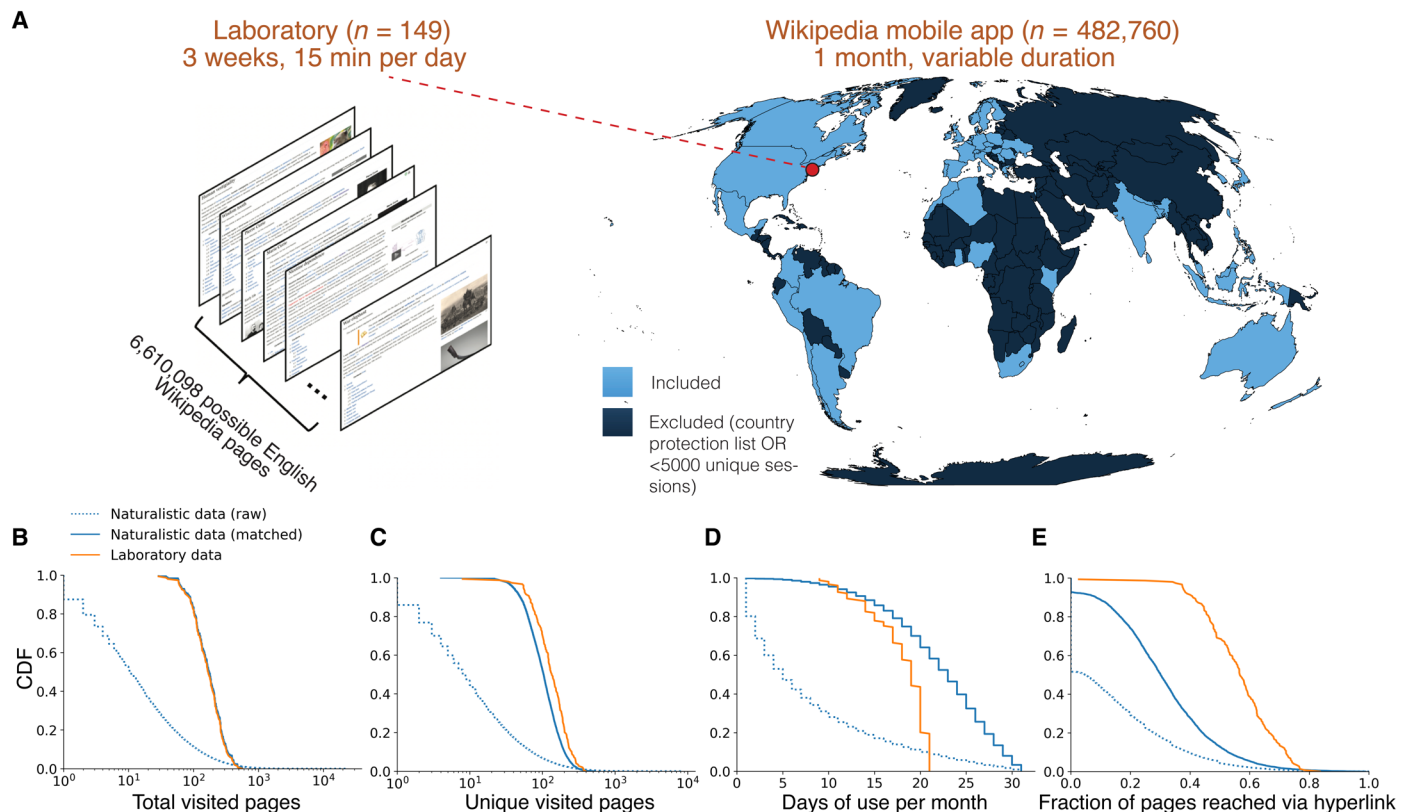


Fig. 1. Geography and summary statistics of laboratory versus naturalistic data. (A) Geography and sample sizes of laboratory and naturalistic data. Naturalistic data includes 14 languages and 50 countries or territories. (B to E) Cumulative distribution functions (CDFs) of summary statistics for laboratory 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).

cultures (49). 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 quantitative evidence for the creative architectural style of the dancer. In making these three contributions, we provide 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.

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 (UPenn) to the naturalistic knowledge networks of readers of Wikipedia across many languages and countries. In the first dataset from UPenn, 149 participants browsed 18,654 unique Wikipedia articles. Participants were instructed to browse English Wikipedia for 15 min 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 (Fig. 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 (Fig. 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 (Fig. 1, B to 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 to fairly compare knowledge networks. Specifically, we use propensity score matching on the number of page views per reader (see Materials and 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 (50). 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 (Fig. 1, B to D). Furthermore, the differences in the distributions of the other summary statistics also become less pronounced, suggesting that differences in the datasets before propensity score matching derive largely from differences in the number of pages viewed per reader (figs. S1 and S2). 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 (36). These metrics measure architectural styles of curiosity that generate tight or loose knowledge networks. Here, we expand this characterization by calculating six additional network metrics (see Materials and 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 (Fig. 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 that have been uncovered in previous studies using the laboratory data generalize to a broader population of Wikipedia's readership in the naturalistic data.

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 Materials and Methods for details). We find that the averaged distance between laboratory and naturalistic data is around $d \approx 0.3$ (Fig. 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 (51, 52). We find that the averaged distance between the naturalistic data and Wikispeedia is $d \geq 0.7$, substantially 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 $d \leq 0.05$, substantially 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.

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 (Fig. 3B). Although for most countries (such as Canada, the United States, 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 Materials and Methods). Comparing the knowledge networks for the naturalistic datasets from different languages (Fig. 3C), we find that distances between English and most languages (German, Russian, Japanese, Spanish, Dutch, Hebrew, Ukrainian, and Chinese) are of comparable magnitude to the distance between the laboratory and naturalistic data ($0.1 \leq d \leq 0.3$). Some languages (Arabic, Hungarian, Bengali, Romanian, and Hindi) show slightly larger distances ($0.3 \leq d \leq 0.4$). However, the latter can be partly

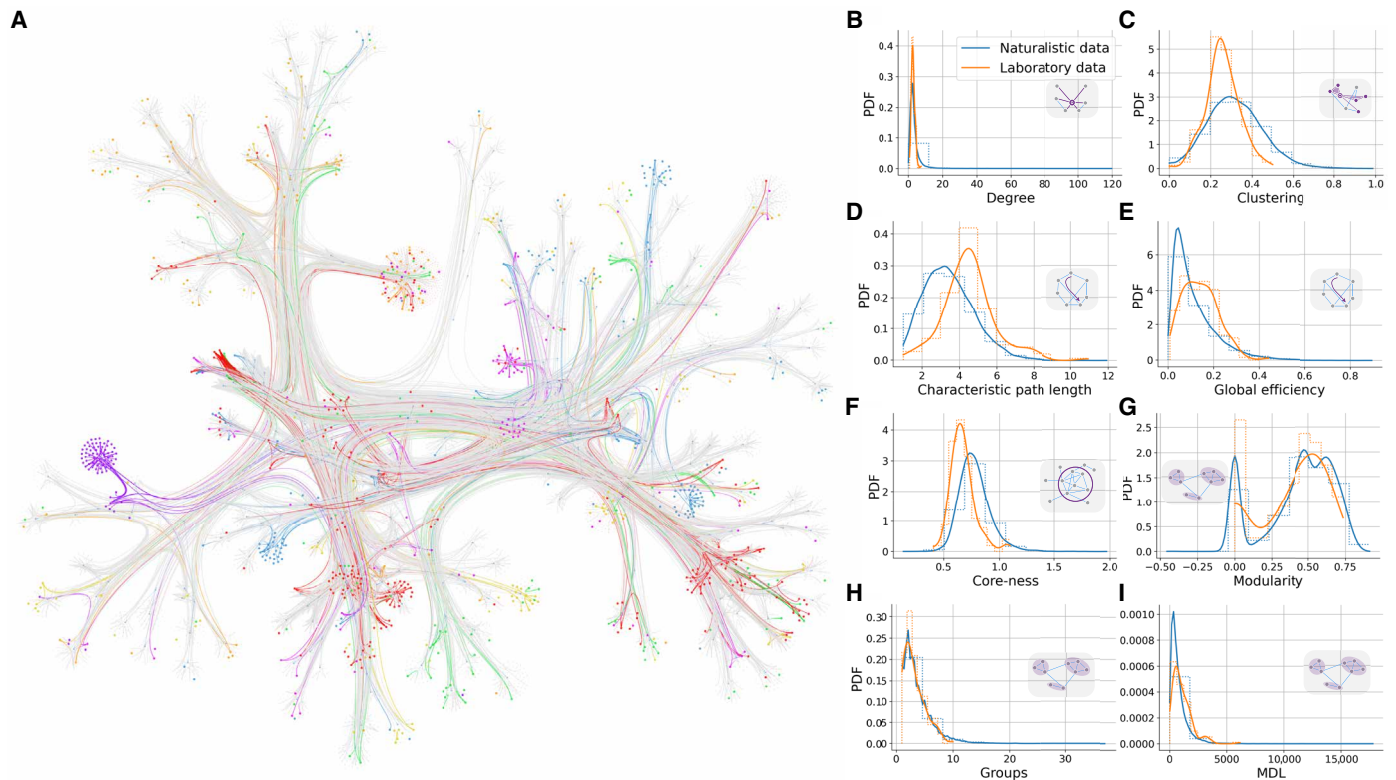


Fig. 2. Distributions of network metrics in knowledge networks from the laboratory and naturalistic data. (A) As readers seek information, they weave diverse temporal threads through nodes defined as Wikipedia articles and across edges defined as hyperlinks. Displayed is a hyperlink network with only 0.1% of nodes in English Wikipedia displayed and different knowledge networks of seven 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 to I) Solid lines indicate the probability density functions (PDFs) 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.

explained by the fact that the underlying networks of hyperlinks between articles are much smaller in size (table S1).

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 Materials and 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 (Fig. 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 (53). 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 ($d \approx 0.3$) (54). The second random walk is biased to undergo an efficient foraging search called a Lévy flight modeled after the laboratory data ($d \approx 0.4$) (36). 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 notable 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

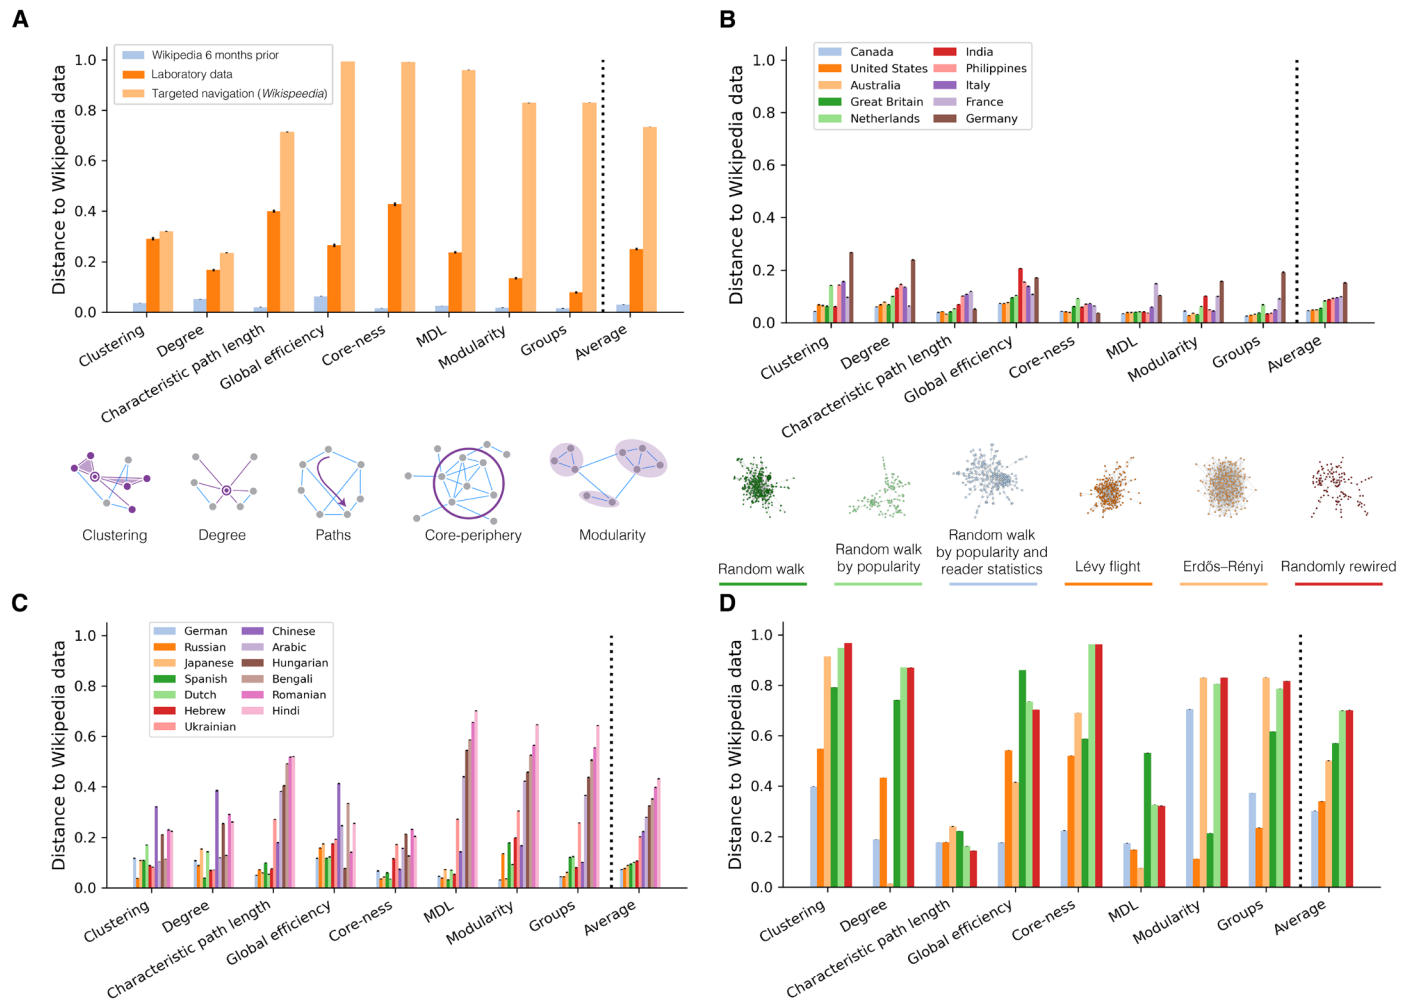


Fig. 3. Distances of the distribution of knowledge network structures between the naturalistic dataset and other datasets. (A) Comparison with three other datasets: a dataset of reading sessions from 2022-03 (instead of 2022-10), a laboratory-based Wikipedia navigation game called Wikispeedia, and laboratory 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.

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 (55–57). 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 (Fig. 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, 36). Specifically, we define a scalar busybody/hunter score by aggregating different network metrics; the score captures whether a knowledge network is more clustered and hunter-like, or more dispersed and busybody-like (Fig. 4A; Materials and 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 ($d_{\text{laboratory}} \approx 0.55, 0.95, 0.86; d_{\text{naturalistic}} \approx 0.59, 0.79, 0.86$) (Fig. 4, B to D). The similar distributions of the busybody/hunter score in the naturalistic and laboratory datasets support the generalizability of the hunter and busybody styles of curiosity.

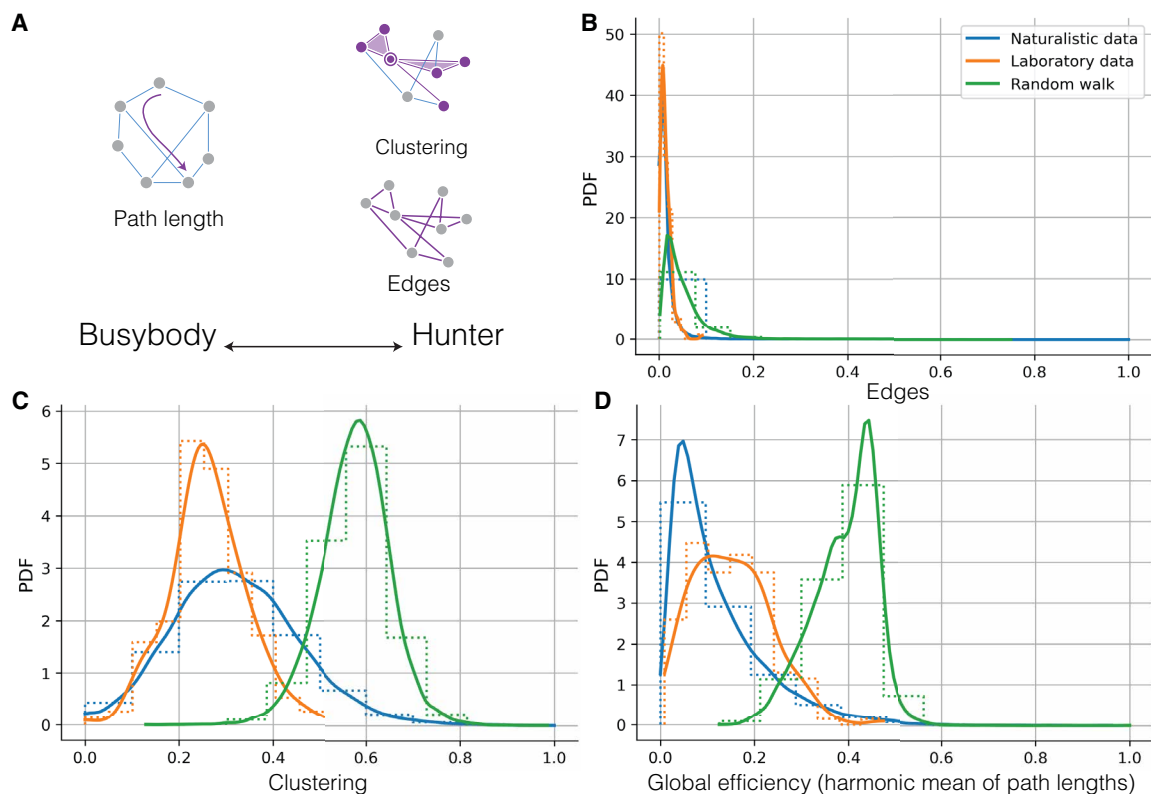


Fig. 4. Similar distributions of busybodies and hunters between the naturalistic and laboratory datasets. (A) The busybody/hunter score is a single metric that characterizes the clustering coefficient and characteristic path lengths in knowledge networks. These metrics were selected because they correlate with deprivation curiosity in the laboratory data (36). 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.

Curiosity styles differ by topic

Readers visit Wikipedia for a variety of reasons and with different information needs (46, 47). Might their curiosity styles differ according to those needs (58)? 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 than hunters (36). 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 that groups articles into a small set of predefined topics (Materials and Methods) (59). For example, high-level topics span four main areas: STEM (science, technology, engineering and math), history and society, geography, and culture (Fig. 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 (Fig. 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 (Fig. 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 and 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 toward 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 (Fig. 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 Materials and Methods for a definition of the information diversity score). The diversity of topics browsed

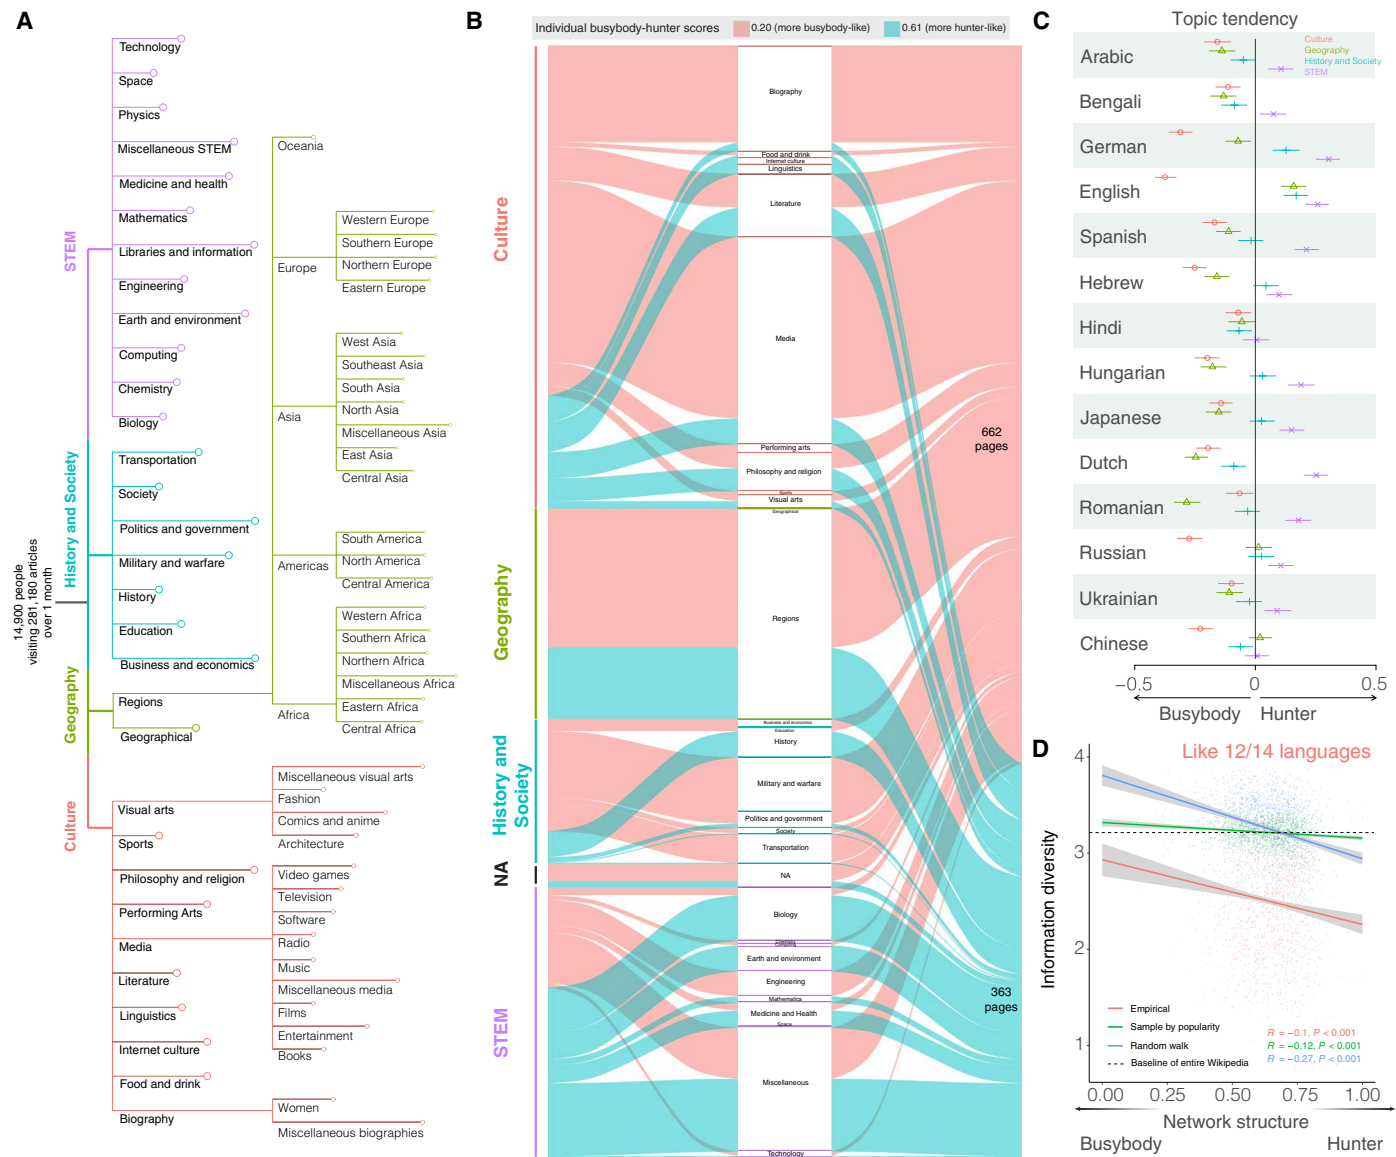


Fig. 5. Curiosity styles evinced by knowledge networks vary by topic. (A) The hierarchy of topics we considered, including 4 broad topics and 64 granular topics. (B) Alluvial diagram that shows divergences among 4 broad topics (left) corresponding to the 64 granular topics (middle) of the articles that were viewed by one hunter-like reader (cyan) and another busybody-like reader (salmon; right). The volume of the streams from left to right indicate the amount of page views that the busybody-like or hunter-like reader distributed to different topics. Starting at the broad topics (left), the culture, geography, and history and society topics were viewed more by the busybody-like reader, whereas the STEM topics were viewed more by the hunter-like reader. Following these streams into the more granular topics (middle), some more specific topics are viewed more by the busybody-like reader such as biography, media, and the performing arts, whereas others were viewed more by the hunter-like reader such as linguistics, biology, and medicine. Last, these streams end in separate pools for each reader, indicating that the busybody-like reader viewed more articles in general than the hunter-like reader. (C) Forest plot of page views (normalized by the page views expected under a random walk) versus the busybody/hunter-score separately for the four 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 are 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.

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.

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 (table S1). 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 exceed 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.

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 (36). 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 (60) and psychological assessments of spatial navigation, a universal requirement across cultures (61). 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 Materials and Methods for details). Geography and ethnic backgrounds can influence the kinds of information sought according to the unique relevance of current events and history (62, 63). Therefore, to investigate this hypothesis, we separated measurements of reader’s information seeking in each country by the different languages of Wikipedia that they accessed. Specifically, we consider measures averaged within datasets of 1490 readers, each from 74 pairs of a given country and a given language edition of Wikipedia (Fig. 6A). To obtain these 74 pairs of a given country and language edition, the aggregated measurements of each country were separated into the 14 different language editions that the reader accessed. This increased granularity is important because our results suggested that language had a stronger

Downloaded from https://www.science.org on December 15, 2024

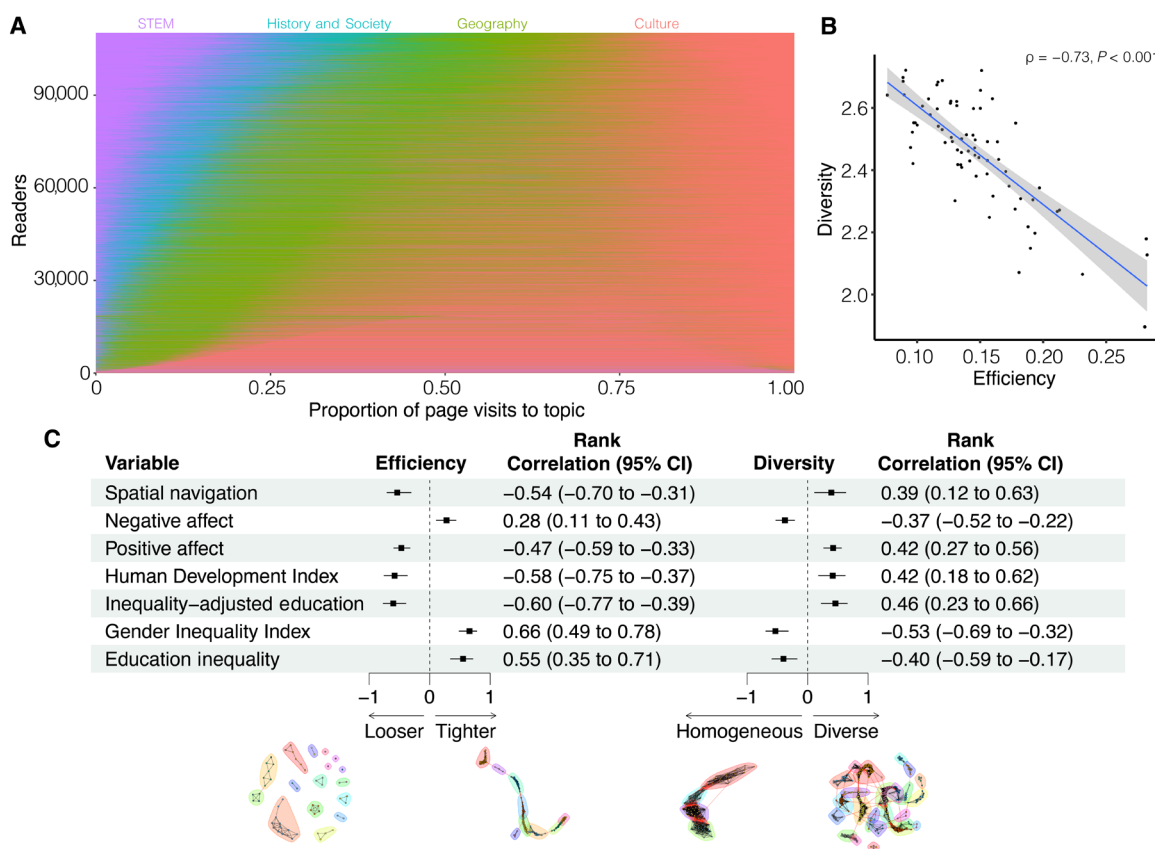


Fig. 6. Diversity and efficiency of knowledge networks predict population-level sociodemographic variables. (A) Readership across 74 pairs of countries and languages vary in the diversity of topics explored, where high diversity (top) corresponds to a more uniform representation of topics compared to low diversity (bottom) with an overrepresentation 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.

effect than geography on knowledge network structure. 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 normalized ratio of the equally distributed level of education to the average observed distribution. Last, we consider gender inequality, which aggregates the maternal mortality ratio, adolescent birth rates, secondary education ratio, and labor force participation ratio. We test whether 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 ($\rho = -0.73$, $P < 0.001$) (Fig. 6B and figs. S3 to S5).

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 ($\rho = -0.55$, $P = 6.74 \times 10^{-6}$) (Fig. 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 ($\rho = -0.40$, $P = 7.26 \times 10^{-3}$) (Fig. 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 highlight the importance of modeling the complex connections coalescing from sequences of information beyond the content of that information. Together, the results are consistent with the notion that certain forms of curiosity support well-being (10, 64) 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 affecting 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 (65). Moreover, that remodeling process has been considered 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 (66–68).

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 (66). 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 (66). 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 redefine edges as the semantic distance between articles (Fig. 7A and Materials and Methods). Consider two simulated walks weaving a temporal thread through five 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 (Fig. 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 ($d \approx 0.78$) (Fig. 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.

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 (Fig. 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. 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.

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 (Fig. 8A). One main axis of this space is the distinction between hunter and busybody scores, from top left to bottom right (Fig. 8A, inset). By including different null models, we can further

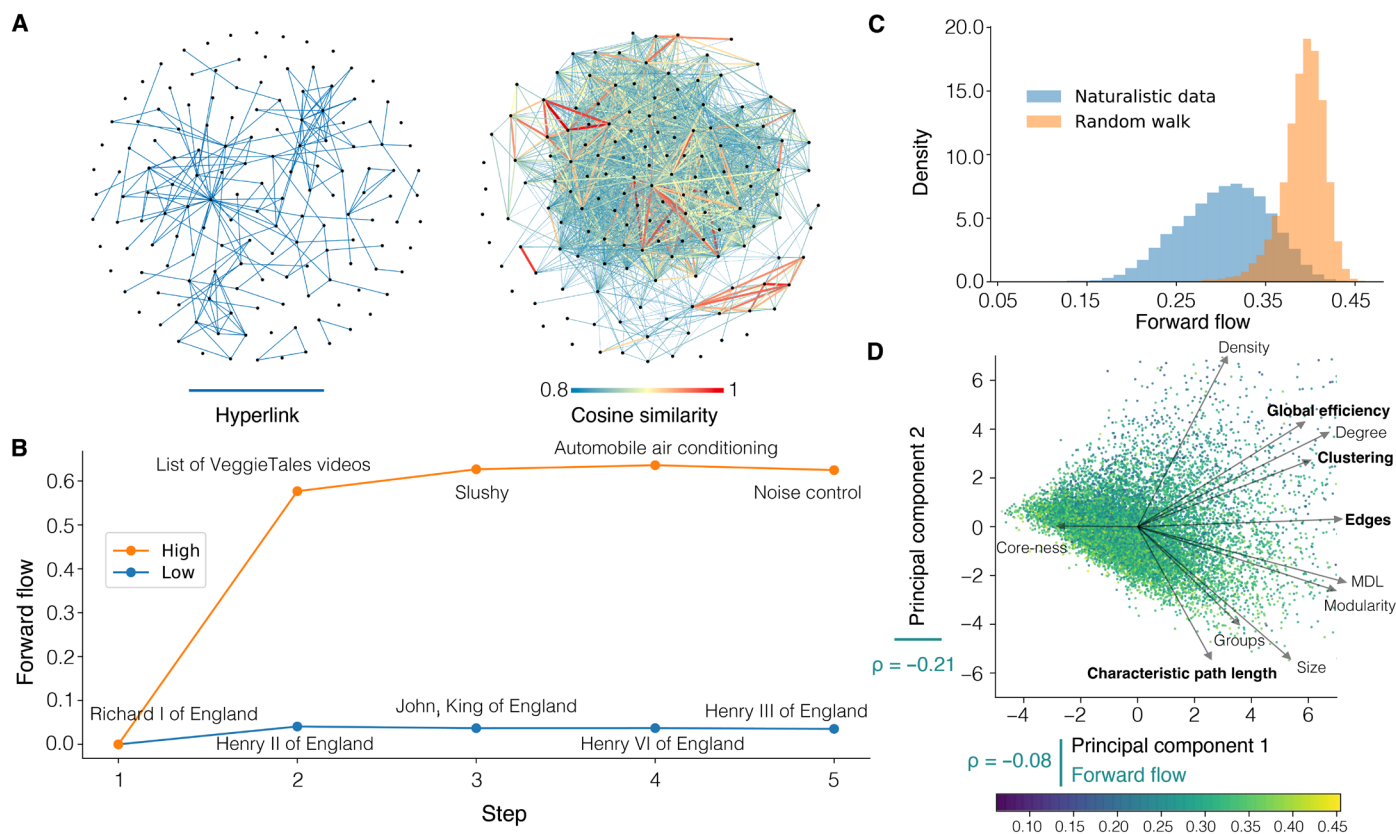


Fig. 7. Operationalizing the dancer style of curiosity in knowledge networks. (A) Forward flow is a topological measure associated with creative thought. The measure is based on a network where edges are defined as the semantic similarity between articles (right) rather than as hyperlinks (left). The semantic network is visualized with edge strengths thresholded above 0.8. (B) Example of the forward flow metric changing during five 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 ($\rho = -0.21$, $P < 0.001$). 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.

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 (Fig. 8B). Knowledge networks from different Wikipedia languages populate similar spaces of curiosity styles. There are some systematic nuances between language versions (Fig. 8C). Some of these differences can be explained by the underlying network structure, such as the different sizes of the hyperlink networks (table S1). 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.

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 laboratory-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 not only that the latter 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 (54).

Our analysis also reveals insights into 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

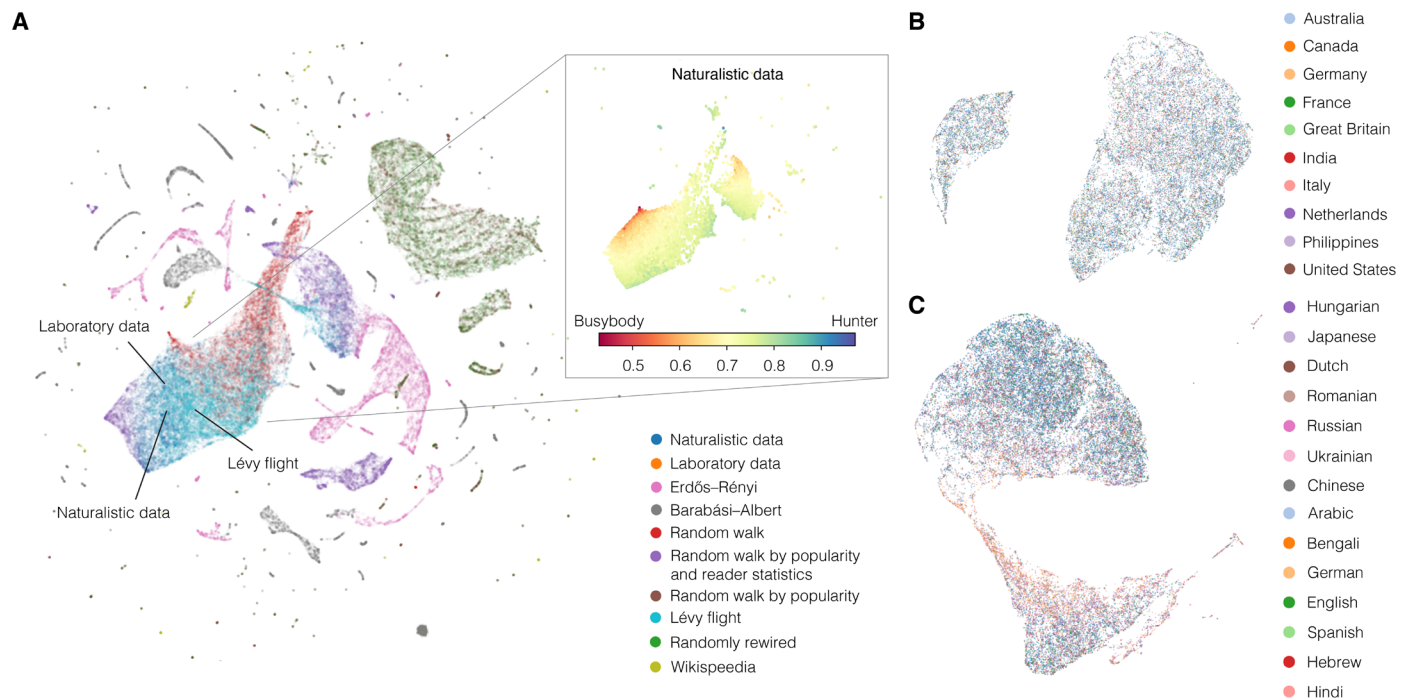


Fig. 8. Uncovering the space of curiosity styles in knowledge networks. (A) All 11 topological metric scores projected onto two dimensions using uniform manifold approximation and projection (UMAP) and colored by dataset or model. Naturalistic data for English Wikipedia are more proximal to the laboratory data and Lévy flight model than other null models (randomly rewired, random walks, Erdős-Rényi, and Barabási-Albert) and browsing tasks (Wikispeedia). Labels mark the centroids of the laboratory, naturalistic, and simulated datasets. Inset: The naturalistic dataset from English Wikipedia is well explained by the overlaid busybody/hunter score. (B and C) Different countries and languages tend to cluster together, suggesting comparable network structure. Two clusters separate due to network size differences (table S1).

readers, partly influenced by their country of access but more substantially by their language.

Curiosity has been increasingly recognized to have a complex association with learning and well-being. For example, curiosity may help allocate attention to information that is neither too complex nor too simple, but to an intermediate “goldilocks” amount of information that is conducive to learning (9, 69–71). Our operationalization of distinct curious practices captures this complexity and underscores the importance of doing so (fig. S6): 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 behavior (72). 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 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 an 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 (73, 74) to form unexpected, remote associations (75). These mechanisms center around the mental maps we create, here operationalized in knowledge networks, by learning the network structure of a sequence of information (76–79).

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 (75, 80). 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 (60). We further tie the architectural styles of information seeking to efficient foraging theories (39, 81–84). 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 (57), positive emotions like interest tend to broaden exploration. Broadened exploration may enable people to build intellectual, emotional, and social resources (57, 85). Animals, including humans, exhibit greater well-being and positive affect with greater experiential diversity (86, 87). Diverse, broad exploration increases novel experiences, which may serve as intrinsic rewards to generate positive affect (88). In support of this idea, the processes of positive affect, foraging, novelty, and exploration may be supported by the integration of activity in the striatum and the hippocampus, brain regions implicated in reward learning, spatial navigation, and novelty detection (86). 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 (42).

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 (89), may relate to inflexibly updating or making rigid the current structures of information (27, 90). 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 (91–93). Such networks may be useful to regulate psychological uncertainty (36) and to help plan especially when the future environment is uncertain (84, 94, 95). The management of uncertainty (9, 71) and the constraints of returning to a central location (or home) versus exploring parts unknown (41, 93, 96) 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, 84, 97). 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 (98), which are based on counting the number of pageviews, can provide a misleading account of the information value that a website provides to readers (99). 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 (27), such insights are relevant in addressing problems of dis- and misinformation in Wikipedia (100) from the perspective of readers. While most of the existing tools and initiatives have been geared toward 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 online (101) using the technique of lateral reading across multiple sources of information (102). 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. Items from the five-dimensional curiosity (5DC) scale (7) overlap with the taxonomy of why readers visit Wikipedia (46). 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 feed) (103) or similarity (related articles shown at the bottom of the mobile site) (104, 105). 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 of knowledge networks is consistent with compression progress theory (42, 106). 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 (41, 78). Seeking information with more connections in learned mental models may support memory (107), perhaps because such connectivity increases the compressibility and navigability of networks (91, 92). 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 (54, 108–112).

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 additional analyses of intensive longitudinal data of day-to-day digital experiences (113, 114). 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 (47). Our finding yields a deeper understanding of how a global readership consumes and connects knowledge across languages and countries. 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 toward knowledge equity (115).

Information seeking can be affected by current events as well as a more stable background concept space consisting of topics with a lasting historical significance (63, 116). Such history and current events can be notably influenced by geography. For example, a large proportion of information seeking behaviors recorded by these studies were on topics involving the United States, the United Kingdom, and Australia likely due to the usage of English Wikipedia. In our laboratory data, we noticed that an individual exhibited strong hunter-like tendencies in seeking information about the past queens of England around the time of Prince Harry's royal wedding. Our work suggests that knowledge network structure differs by language and geography, consistent with different use cases of Wikipedia associated with different socio-economic and cultural backgrounds (47), yet can be analyzed integratively along busybody-hunter and dancer dimensions.

It is worthwhile to consider the associations between hunter-like curiosity and both gender and education inequality. Differences in education and gender are related to information seeking in Wikipedia as well as curiosity-driven information seeking (117–121). Such differences can be investigated in light of prior work in critical social theory (122), intersectional concerns of epistemic freedoms and knowledge production (123), and links between digital inequality and other forms of social inequality including those encompassing education and gender (124). 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 (125, 126) 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. Men tend to have greater specific curiosity than diversive curiosity, where a specific curiosity involves seeking detailed information about a specific topic, whereas diversive curiosity involves seeking a broad range of new information (119).

A second possible explanation relates to gender disparities in readership: It is possible that in countries characterized by greater gender and education inequality, the readers of Wikipedia tend to be people who are more able to access digital resources, are socialized as men, and trained to evince hunter-like curiosity over busybody-like or dancer-like curiosity. In support of this notion, in studies of trait curiosity, men tend to evince specific (hunter-like) versus diversive (busybody-like) styles of epistemic curiosity for obtaining knowledge about objects, things, and the physical sciences (119, 120). Further, people with disadvantaged backgrounds encounter compounding barriers that reinforce their digital exclusion (124). Focusing specifically on Wikipedia readers, it is relevant to note that the level of resources for education in different countries influences the topics, duration, and depth of information seeking (47, 99), and women historically have had fewer of those resources. Global online surveys have shown that women are underrepresented among readers of Wikipedia (117).

A third possible explanation relates to epistemic norms in different disciplines or areas of knowledge. Intrinsic interest or lack thereof in topics such as mathematics, physics, and engineering is

sometimes used to explain skewed gender proportions in different fields and levels of STEM education (127, 128), alongside gendered stereotypes and role models unique to a culture and geography (129). Our present results suggest a relationship between information seeking about STEM and hunter-like curiosity, which is analogous to the kind of specific versus diversive epistemic curiosity found to be more prevalent in men (119). 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 three candidate drivers, it would be interesting to test whether action to support diverse curiosity styles might effectively work against hegemonic norms, as people with high curiosity tend to exhibit flexibility with gender norms, rebelliousness, unconventionality, and social perceptiveness (130). However, it is also important to acknowledge that readers in locations of marked social injustice (including gender and education inequality) could effectively use any sort of curiosity (including hunter-like curiosity) to build resistant knowledge (17, 131). In line with this idea, the key knowledge network statistic might be one that is dynamic, such as state entropy, to characterize how flexibly someone moves between diverse modes of curiosity (41, 42).

Last, 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” (48). Our work on the curiosity of Wikipedia's readership constitutes an example of how this laboratory contributes to addressing current challenges in human behavior research—not only by increasing ecological validity or diversifying representation empirically (132) but also by advancing theories of human behavior (133). More generally, Wikipedia plays a central role in contributing to the development and improvement of the scientific understanding of meaningful measures of human society (134); 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 behavior on social media sites (135).

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 (136). Future work could also assess the effect of external events, such as news related to current events, across multiple languages and countries. Knowledge network structure was relatively similar despite being separated by 6 months in the naturalistic dataset. This may suggest that individual information seeking styles are fairly consistent across time, characterizing a trait-like tendency related to a curiosity component like deprivation sensitivity. Because we used a cross-sectional analysis across time points, this similarity could also suggest that the prevalence of different information seeking styles is robust across time. To test the stability of long-term trait-like styles or the transience of state-like styles of information seeking, future research could collect a longitudinal sample across several months. Prior work studying the variability in hunter and busybody styles over time indicated that while it is meaningful to talk about trait differences in knowledge network

building, there are also within-person changes in building tendencies across time (36). Clustering methods such as a hidden Markov model could be used to identify three styles, characterizing hunters, busybodies, and dancers, or more. We expect that beyond a trait-like preference to dwell in certain styles, individuals flexibly shift between modes of curiosity perhaps beyond the three we have examined here.

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 (54), 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 (137, 138); for example, redundant subindices 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 datasets and aggregate measures. Third, we lacked data on individual characteristics of Wikipedia readers in the naturalistic dataset. Therefore we used, as a proxy, separate survey data on mood and well-being collected from a population sample across many countries (139). The present work builds upon prior research that noted that geographical characteristics may explain differences in the patterns and topics of information seeking (47, 63). This approach has previously been used to link individual motivations and goals of information seeking on Wikipedia to account for the differing socio-economic or cultural background of readers (47). This method allows us to offer preliminary evidence of a consistent relationship between mood and information seeking at both an individual (laboratory data) and population level (aggregated naturalistic data). However, aggregated measurements at the level of countries obscure unique variability among individuals, known to influence individual information seeking behavior (62, 116, 140). The validity of this aggregate analysis hinges on how representative the naturalistic data and survey methods are to the population, which can be affected by sampling biases.

MATERIALS AND 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 (31, 36, 64). 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, SD = 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%), and 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 1830 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 5 min to complete. These surveys included scales for depressed mood and anxiety (141).

The second component came immediately after the daily diary and was a 15-min 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 min 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://wikipedia.org/>. We would like you to spend 15 min 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 min 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 were collected for the months of March 2022 and October 2022 from the webrequest logs (142). In general, these data are 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 bots (143). We identify unique readers via the *wm-fuuid* (144) associated with the app installation (readers of the app have the option to turn off sending this ID with their requests). 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 list (145); dropping IP, user agent, and fine-grained geospatial information; and 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* (146). 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 1490 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, 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 (47).
- 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 5000 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.

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 (51, 52). 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 4604 articles. The task

was successfully completed by 14,246 participants and took an average of 2.64 ± 6.01 min. The data were downloaded from an <https://snap.stanford.edu/data/wikispeedia.html> 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 (39, 73). To index spatial navigation abilities, we evaluated previously collected data on nonverbal spatial navigation using a wayfinding task in 397,162 participants from 38 countries (147). 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 (148, 149). Performance is also influenced by one's environment (61), 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 country-wide indicators from the Human Development Index (2021), Gender Inequality Index (2021), and World Happiness Report (2021–2022) (139, 150). 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 https://hdr.undp.org/sites/default/files/2021-22_HDR/hdr2021-22_technical_notes.pdf 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 table (151) 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, core-ness, 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 (152–154).

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

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

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.

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

$$c_u = \frac{2T(n)}{k(n)[k(n)-1]} \tag{2}$$

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

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

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

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

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 (155). The core-ness can be defined as

$$Q_C = \frac{1}{v_C} \left[\sum_{ij \in C_c} (w_{ij} - \gamma_C \bar{w}) - \sum_{ij \in C_p} (w_{ij} - \gamma_C \bar{w}) \right] \tag{5}$$

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 (156). 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

$$P(b|A) = \frac{P(A|\theta,b)P(\theta,b)}{P(A)} \tag{6}$$

where A is the network's adjacency matrix and θ 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

$$Q = \frac{1}{2M} \sum_r m_{rr} - \frac{m_r^2}{2M} \tag{7}$$

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

Downloaded from https://www.science.org on December 15, 2024

edges. Busybodies tend to construct looser networks with weaker within-module connections and hence less modularity than hunters.

The minimum description length (157) is the amount of information required to describe the data using our fitted stochastic block model. Maximizing the posterior probability of $P(b | A)$ gives the minimum description length Σ defined as

$$\Sigma = -\ln P(A | \theta, b) - \ln P(\theta, b) \quad (8)$$

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 (36). 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 (66). 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 have 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' pretrained word-vector embeddings from fastText (158).

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 (54), and a foraging model rooted in theories from ecology and psychology (84, 159).

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 (160). The Erdős-Rényi network was generated using an edge probability corresponding to the edge density of the reader's network size (161). The Barabási-Albert model generates scale-free networks with heterogeneous degree distribution by a preferential attachment (rich-get-richer) mechanism (162).

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 was 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 (163). 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 (53).

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 (164). 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 (54).

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 (36, 41). Here, we repeatedly draw pairs of parameters with replacement from the empirical distribution 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, 43), 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 (81). 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 (159). 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 (36).

Information diversity

The diversity of information is important for conceptualizations of epistemic curiosity. This kind of curiosity is a drive for knowledge. In a historic-philosophical framework of curiosity styles, there are two complementary inclinations to seek knowledge. Hunters tend to pursue a set of similar topics, whereas busybodies tend to explore a wide range of 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 STEM (59). 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. Last, 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- n matrices of network structure (165). 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. Last, we used the two-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 (166–174). 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 (170, 175). 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 (i) names, pronouns, and social media profiles used to construct the databases may not, in every case, be indicative of gender identity and (ii) it cannot account for intersex, nonbinary, 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 (176, 177). 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 (i) names and Florida Voter Data to make the predictions may not be indicative of racial/ethnic identity, and (ii) 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.

Supplementary Materials

This PDF file includes:

Supplementary Materials and Methods
Supplementary Text
Figs. S1 to S6
Table S1

REFERENCES AND NOTES

- G. A. Miller, in *The Study of Information: Interdisciplinary Messages*, F. Machlup, U. Mansfield, Eds. (Wiley, 1983), pp. 111–113.
- P. Pirolli, S. Card, Information foraging. *Psychol. Rev.* **106**, 643–675 (1999).
- D. E. Berlyne, Curiosity and exploration. *Science* **153**, 25–33 (1966).
- D. E. Berlyne, A theory of human curiosity. *Br. J. Psychol.* **45**, 180–191 (1954).
- P. J. Silvia, Interest—The curious emotion. *Curr. Dir. Psychol. Sci.* **17**, 57–60 (2008).
- C. Kidd, B. Y. Hayden, The psychology and neuroscience of curiosity. *Neuron* **88**, 449–460 (2015).
- T. B. Kashdan, D. J. Disabato, F. R. Goodman, P. E. McKnight, The Five-Dimensional Curiosity Scale Revised (5DCR): Briefer subscales while separating overt and covert social curiosity. *Pers. Individ. Dif.* **157**, 109836 (2020).
- J. Gottlieb, P.-Y. Oudeyer, M. Lopes, A. Baranes, Information seeking, curiosity and attention: Computational and neural mechanisms. *Trends Cogn. Sci.* **17**, 585–593 (2013).
- J. Gottlieb, P.-Y. Oudeyer, Towards a neuroscience of active sampling and curiosity. *Nat. Rev. Neurosci.* **19**, 758–770 (2018).
- T. B. Kashdan, M. C. S., D. J. D., P. E. M. K., J. Bekier, J. Kaji, R. Lazarus, The five-dimensional curiosity scale: Capturing the bandwidth of curiosity and identifying four unique subgroups of curious people. *J. Res. Pers.* **73**, 130–149 (2018).
- P. Zurn, in *Toward New Philosophical Explorations of the Epistemic Desire to Know: Just Curious about Curiosity* (Cambridge Scholars Publishing, 2019), pp. 27–49.
- P. Zurn, D. S. Bassett, *Curious Minds: The Power of Connection* (MIT Press, 2022).
- D. E. Blasi, J. Henrich, E. Adamou, D. Kemmerer, A. Majid, Over-reliance on English hinders cognitive science. *Trends Cogn. Sci.* **26**, 1153–1170 (2022).
- J. Henrich, S. J. Heine, A. Norenzayan, The weirdest people in the world? *Behav. Brain Sci.* **33**, 61–83 (2010).
- P. Zurn, D. S. Bassett, On curiosity: A fundamental aspect of personality, a practice of network growth. *Personal. Neurosci.* **1**, e13 (2018).
- J. Medina, *The Epistemology of Resistance: Gender and Racial Oppression, Epistemic Injustice, and Resistant Imaginations* (Oxford Univ. Press, 2012).
- P. Zurn, *Curiosity and Power: The Politics of Inquiry* (University of Minnesota Press, 2021).
- P. Zurn, Curiosity: An affect of resistance. *Theory Event* **24**, 611–617 (2021).
- N. Park, C. Peterson, M. E. Seligman, Strengths of character and well-being. *J. Soc. Clin. Psychol.* **23**, 603 (2004).
- M. J. Kang, M. Hsu, I. M. Krajbich, G. Loewenstein, S. M. McClure, J. T. Y. Wang, C. F. Camerer, The wick in the candle of learning: Epistemic curiosity activates reward circuitry and enhances memory. *Psychol. Sci.* **20**, 963–973 (2009).
- M. J. Gruber, B. D. Gelman, C. Ranganath, States of curiosity modulate hippocampus-dependent learning via the dopaminergic circuit. *Neuron* **84**, 486–496 (2014).
- S. Wade, C. Kidd, The role of prior knowledge and curiosity in learning. *Psychon. Bull. Rev.* **26**, 1377–1387 (2019).
- P. Freire, in *Toward a Sociology of Education* (Routledge, 2020), pp. 374–386.
- J. Clark, A. Vincent, X. Wang, A. L. McGowan, D. M. Lydon-Staley, Curiosity, surprise, and the recall of tobacco-related health information in adolescents. *J. Health Commun.* **28**, 446–457 (2022).
- T. Lyew, A. Ikhtlas, F. Sayed, A. Vincent, D. Lydon-Staley, Curiosity, surprise, and the recall of tobacco-related health information in adolescents. *J. Health Commun.* **28**, 446–457 (2023).
- A. Schumacher, Y. Kammerer, C. Scharinger, S. Gottschling, N. Hübner, M. Tibus, How do Intellectually Curious and Interested People Learn and Attain Knowledge? A Focus on Behavioral Traces of Information Seeking (2024); <https://doi.org/10.31219/osf.io/6djkr>.
- C. M. Zedelius, M. E. Gross, J. W. Schooler, Inquisitive but not discerning: Deprivation curiosity is associated with excessive openness to inaccurate information. *J. Res. Pers.* **98**, 104227 (2022).
- U. Schiefele, Interest and learning from text. *Sci. Stud. Read.* **3**, 257–259 (1999).
- B. Renner, Curiosity about people: The development of a social curiosity measure in adults. *J. Pers. Assess.* **87**, 305–316 (2006).
- L. D. Kaczmarek, B. Bączkowski, J. Enko, B. Baran, P. Theuns, Subjective well-being as a mediator for curiosity and depression. *Psychol. Bull.* **45**, 200–204 (2014).
- D. M. Lydon-Staley, P. Zurn, D. S. Bassett, Within-person variability in curiosity during daily life and associations with well-being. *J. Pers.* **88**, 625–641 (2020).
- P. J. Silvia, A. P. Christensen, Looking up at the curious personality: Individual differences in curiosity and openness to experience. *Curr. Opin. Behav. Sci.* **35**, 1–6 (2020).
- J. Kristeva, *Black Sun: Depression and Melancholia* (Columbia Univ. Press, 1989).
- A. Steglich-Petersen, S. Varga, Curiosity and zetetic style in ADHD. *Philos. Psychol.* **1–25** (2023).
- K. P. Lindgren, P. M. Mullins, C. Neighbors, J. A. Blayney, Curiosity killed the cocktail? Curiosity, sensation seeking, and alcohol-related problems in college women. *Addict. Behav.* **35**, 513–516 (2010).
- D. M. Lydon-Staley, D. Zhou, A. S. Blevins, P. Zurn, D. S. Bassett, Hunters, busybodies and the knowledge network building associated with deprivation curiosity. *Nat. Hum. Behav.* **5**, 327–336 (2021).
- G. Loewenstein, The psychology of curiosity: A review and reinterpretation. *Psychol. Bull.* **116**, 75–98 (1994).
- T. Li, H. Huang, J. Liu, X. Tang, Killing the cats or satisfying the human? The role of epistemic curiosity in adolescents' multidimensional well-being. *J. Pac. Rim Psychol.* **17**, 10.1177/18344909231185381 (2023).
- P. M. Todd, T. T. Hills, Foraging in MIND. *Curr. Dir. Psychol. Sci.* **29**, 309–315 (2020).
- D. S. Bassett, in *Curiosity Studies: A New Ecology of Knowledge*, P. Zurn, A. Shankar, Eds. (University of Minnesota Press, 2020).
- D. Zhou, D. M. Lydon-Staley, P. Zurn, D. S. Bassett, The growth and form of knowledge networks by kinesthetic curiosity. *Curr. Opin. Behav. Sci.* **35**, 125–134 (2020).
- S. P. Patankar, D. Zhou, C. W. Lynn, J. Z. Kim, M. Ouellet, H. Ju, P. Zurn, D. M. Lydon-Staley, D. S. Bassett, Curiosity as filling, compressing, and reconfiguring knowledge networks. [arXiv:2204.01182 \[q-bio.NC\]](https://arxiv.org/abs/2204.01182) (2022).
- J. A. Litman, T. L. Jimerson, The measurement of curiosity as a feeling of deprivation. *J. Pers. Assess.* **82**, 147–157 (2004).
- Wikistats—Statistics for wikimedia projects. <https://stats.wikimedia.org/>.
- G. Domantas, The most visited website in every country (that isn't a search engine) (2022); <https://hostinger.com/tutorials/the-most-visited-website-in-every-country>.
- P. Singer, F. Lemmerich, R. West, L. Zia, E. Wulczyn, M. Strohmaier, J. Leskovec, in *Proceedings of the 26th International Conference on World Wide Web* (Association for Computing Machinery, 2017), pp. 1591–1600.
- F. Lemmerich, D. Sáez-Trumper, R. West, L. Zia, in *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining* (Association for Computing Machinery, 2019), pp. 618–626.
- B. M. Hill, A. Shaw, in *Wikipedia @ 20* (The MIT Press, 2020), pp. 159–174.
- M. Gerlach, Research: Understanding curious and critical readers. <https://meta.wikimedia.org/wiki/Research:Understanding-Curious-and-Critical-Readers>.
- P. R. Rosenbaum, D. B. Rubin, The central role of the propensity score in observational studies for causal effects. *Biometrika* **70**, 41–55 (1983).
- R. West, J. Pineau, D. Precup, *Twenty-First International Joint Conference on Artificial Intelligence* (Association for Computing Machinery, 2009).
- R. West, J. Leskovec, in *Proceedings of the 21st International Conference on World Wide Web* (Association for Computing Machinery, 2012), pp. 619–628.
- J. T. Abbott, J. L. Austerweil, T. L. Griffiths, in *Neural Information Processing Systems Conference; A Preliminary Version of This Work Was Presented at the Aforementioned Conference* (American Psychological Association, 2015), vol. 122, p. 558.
- T. Piccardi, M. Gerlach, A. Arora, R. West, A large-scale characterization of how readers browse wikipedia. *ACM Transac. Web* **17**, 1–22 (2023).
- A. S. Keller, J. E. Leikauf, B. Holt-Gosselin, B. R. Staveland, L. M. Williams, Paying attention to attention in depression. *Transl. Psychiatry* **9**, 279 (2019).
- A. J. Whitmer, I. H. Gotlib, An attentional scope model of rumination. *Psychol. Bull.* **139**, 1036–1061 (2013).
- B. L. Fredrickson, The role of positive emotions in positive psychology. *Am. Psychol.* **56**, 218–226 (2001).
- C. A. Kelly, T. Sharot, Individual differences in information-seeking. *Nat. Commun.* **12**, 7062 (2021).
- I. Johnson, M. Gerlach, D. Sáez-Trumper, in *Companion Proceedings of the Web Conference 2021* (Association for Computing Machinery, 2021), pp. 594–601.
- H. van Praag, G. Kempermann, F. H. Gage, Neural consequences of environmental enrichment. *Nat. Rev. Neurosci.* **1**, 191–198 (2000).
- A. Coutrot, R. Silva, E. Manley, W. de Cothi, S. Sami, V. D. Bohbot, J. M. Wiener, C. Hölcher, R. C. Dalton, M. Hornberger, H. J. Spiers, Global determinants of navigation ability. *Curr. Biol.* **28**, 2861–2866.e4 (2018).
- M. Zhu, T. Yasserli, J. Kertész, Individual differences in knowledge network navigation. *Sci. Rep.* **14**, 8331 (2024).

63. P. Gildersleve, R. Lambiotte, T. Yasseri, Between news and history: Identifying networked topics of collective attention on Wikipedia. *J. Comput. Soc. Sci.* **6**, 845–875 (2023).
64. D. M. Lydon-Staley, E. B. Falk, D. S. Bassett, Within-person variability in sensation-seeking during daily life: Positive associations with alcohol use and self-defined risky behaviors. *Psychol. Addict. Behav.* **34**, 257–268 (2020).
65. B. Tversky, K. Hemenway, Objects, parts, and categories. *J. Exp. Psychol.* **113**, 169 (1984).
66. K. Gray, S. Anderson, E. E. Chen, J. M. Kelly, M. S. Christian, J. Patrick, L. Huang, Y. N. Kenett, K. Lewis, "Forward flow": A new measure to quantify free thought and predict creativity. *Am. Psychol.* **74**, 539–554 (2019).
67. Y. N. Kenett, D. Anaki, M. Faust, Investigating the structure of semantic networks in low and high creative persons. *Front. Hum. Neurosci.* **8**, 407 (2014).
68. Y. N. Kenett, M. Faust, A semantic network cartography of the creative mind. *Trends Cogn. Sci.* **23**, 271 (2019).
69. C. Kidd, S. T. Piantadosi, R. N. Aslin, The goldilocks effect: Human infants allocate attention to visual sequences that are neither too simple nor too complex. *PLOS ONE* **7**, e36399 (2012).
70. P.-Y. Oudeyer, L. B. Smith, How evolution may work through curiosity-driven developmental process. *Top. Cogn. Sci.* **8**, 492–502 (2016).
71. R. Dubey, T. L. Griffiths, Reconciling novelty and complexity through a rational analysis of curiosity. *Psychol. Rev.* **127**, 455–476 (2020).
72. Replications do not fail. *Nat. Hum. Behav.* **4**, 559 (2020).
73. J. L. Bellmund, P. Gärdenfors, E. I. Moser, C. F. Doeller, Navigating cognition: Spatial codes for human thinking. *Science* **362**, eaat6766 (2018).
74. J. Aru, M. Drüke, J. Pikamäe, M. E. Larkum, Mental navigation and the neural mechanisms of insight. *Trends Neurosci.* **46**, 100–109 (2023).
75. C. B. Horton, M. F. Mason, Getting curiouser and curiouser about creativity: The search for a nuanced model. *Behav. Brain Sci.* **47**, e102 (2024).
76. D. Schiller, H. Eichenbaum, E. A. Buffalo, L. Davachi, D. J. Foster, S. Leutgeb, C. Ranganath, Memory and space: Towards an understanding of the cognitive map. *J. Neurosci.* **35**, 13904–13911 (2015).
77. R. A. Epstein, E. Z. Patai, J. B. Julian, H. J. Spiers, The cognitive map in humans: Spatial navigation and beyond. *Nat. Neurosci.* **20**, 1504–1513 (2017).
78. I. Momennejad, Learning structures: Predictive representations, replay, and generalization. *Curr. Opin. Behav. Sci.* **32**, 155–166 (2020).
79. N. W. Morton, A. R. Preston, Concept formation as a computational cognitive process. *Curr. Opin. Behav. Sci.* **38**, 83–89 (2021).
80. K. Christoff, Z. C. Irving, K. C. Fox, R. N. Spreng, J. R. Andrews-Hanna, Mind-wandering as spontaneous thought: A dynamic framework. *Nat. Rev. Neurosci.* **17**, 718–731 (2016).
81. G. M. Viswanathan, M. G. Da Luz, E. P. Raposo, H. E. Stanley, *The Physics of Foraging: An Introduction to Random Searches and Biological Encounters* (Cambridge Univ. Press, 2012).
82. T. Rhodes, M. T. Turvey, Human memory retrieval as Lévy foraging. *Phys. A Stat. Mech. Appl.* **385**, 255–260 (2007).
83. J. Zhu, A. Sanborn, N. Chater, Mental sampling in multimodal representations, in *Advances in Neural Information Processing Systems* (MIT Press, 2018), pp. 5748–5759.
84. K. Garg, C. T. Kello, Efficient Lévy walks in virtual human foraging. *Sci. Rep.* **11**, 5242 (2021).
85. B. L. Fredrickson, in *Advances in Experimental Social Psychology* (Elsevier, 2013), vol. 47, pp. 1–53.
86. A. S. Heller, T. C. Shi, C. E. C. Ezie, T. R. Reneau, L. M. Baez, C. J. Gibbons, C. A. Hartley, Association between real-world experiential diversity and positive affect relates to hippocampal–striatal functional connectivity. *Nat. Neurosci.* **23**, 800–804 (2020).
87. N. M. Saragosa-Harris, A. O. Cohen, T. R. Reneau, W. J. Villano, A. S. Heller, C. A. Hartley, Real-world exploration increases across adolescence and relates to affect, risk taking, and social connectivity. *Psychol. Sci.* **33**, 1664–1679 (2022).
88. D. E. Berlyne, Novelty, complexity, and hedonic value. *Percept. Psychophys.* **8**, 279–286 (1970).
89. A. Hsiung, J.-H. Poh, S. A. Huettel, R. A. Adcock, Curiosity evolves as information unfolds. *Proc. Natl. Acad. Sci. U.S.A.* **120**, e2301974120 (2023).
90. A. W. Kruglanski, D. M. Webster, in *The Motivated Mind* (Taylor and Francis, 2018), pp. 60–103.
91. C. W. Lynn, D. S. Bassett, Quantifying the compressibility of complex networks. *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2023473118 (2021).
92. M. Boguna, D. Krioukov, K. C. Claffy, Navigability of complex networks. *Nat. Phys.* **5**, 74–80 (2009).
93. K. Garg, C. Padilla-Iglesias, N. Restrepo Ochoa, V. B. Knight, Hunter–gatherer foraging networks promote information transmission. *R. Soc. Open Sci.* **8**, 211324 (2021).
94. F. Bartumeus, F. Peters, S. Pueyo, C. Marrasé, J. Catalan, Helical Lévy walks: Adjusting searching statistics to resource availability in microzooplankton. *Proc. Natl. Acad. Sci. U.S.A.* **100**, 12771–12775 (2003).
95. D. Zhou, A. M. Bornstein, Expanding horizons in reinforcement learning for curious exploration and creative planning. *Behav. Brain Sci.* **47**, e118 (2024).
96. A. Gopnik, Childhood as a solution to explore–exploit tensions. *Philos. Trans. R. Soc. B* **375**, 20190502 (2020).
97. N. C. Harhen, A. M. Bornstein, Overharvesting in human patch foraging reflects rational structure learning and adaptive planning. *Proc. Natl. Acad. Sci. U.S.A.* **120**, e2216524120 (2023).
98. L. Hong, M. Lalmas, in *Companion Proceedings of the 2019 World Wide Web Conference* (Association for Computing Machinery, 2019), pp. 1303–1305.
99. N. TeBlunthuis, T. Bayer, O. Vasileva, in *Proceedings of the 15th International Symposium on Open Collaboration* (Association for Computing Machinery, 2019), pp. 1–14.
100. D. Sáez-Trumper, Disinformation and AI: The differences between wikipedia and social media (2021).
101. J. Green, Crash course video #5: Using wikipedia (2019).
102. S. Wineburg, S. McGrew, Lateral reading and the nature of expertise: Reading less and learning more when evaluating digital information. *Teach. Coll. Rec.* **121**, 1–40 (2019).
103. Wikimedia apps/feeds. MediaWiki. https://www.mediawiki.org/wiki/Wikimedia_Apps/Feeds (2024).
104. E. Wulczyn, Wikipedia Navigation Vectors. figshare. Dataset (2017); <https://doi.org/10.6084/m9.figshare.3146878.v6>.
105. Extension: Related articles. <https://mediawiki.org/wiki/Extension:RelatedArticles>.
106. S. P. Patankar, M. Ouellet, J. Cervino, A. Ribeiro, K. A. Murphy, D. S. Bassett, Intrinsically motivated graph exploration using network theories of human curiosity in *Proceedings of the Second Learning on Graphs Conference* PMLR 231 (2024).
107. H. Lee, J. Chen, Predicting memory from the network structure of naturalistic events. *Nat. Commun.* **13**, 4235 (2022).
108. D. Lamprecht, K. Lerman, D. Helic, M. Strohmaier, How the structure of Wikipedia articles influences user navigation. *New Rev. Hypermedia Multimed.* **23**, 29–50 (2017).
109. G. C. Rodi, V. Loreto, F. Tria, Search strategies of Wikipedia readers. *PLOS ONE* **12**, e0170746 (2017).
110. K. Ruggeri, B. Večkalov, L. Bojanić, T. L. Andersen, S. Ashcroft-Jones, N. Ayacaxli, P. Barea-Arroyo, M. L. Berge, L. D. Bjørndal, A. Bursaloğlu, V. Bühler, M. Čadek, M. Çetinçelik, G. Clay, A. Cortijos-Bernabeu, K. Damjanović, T. M. Dugue, M. Esberg, C. Esteban-Serna, E. N. Felder, M. Friedemann, D. I. Frontera-Villanueva, P. Gale, E. Garcia-Garzon, S. J. Geiger, L. George, A. Girardello, A. Gracheva, A. Gracheva, M. Guillory, M. Hecht, K. Herte, B. Hubená, W. Ingalls, L. Jakob, M. Janssens, H. Jarke, O. Kácha, K. N. Kalinova, R. Karakasheva, P. R. Khorrami, Ž. Lep, S. Lins, I. S. Lofthus, S. Mamede, S. Mareva, M. F. Mascarenhas, L. McGill, S. Morales-Izquierdo, B. Moltrecht, T. S. Mueller, M. Musetti, J. Nelsson, T. Otto, A. F. Paul, I. Pavlović, M. B. Petrović, D. Popović, G. M. Prinz, J. Razum, I. Sakelarijev, V. Samuels, I. Sanguino, N. Say, J. Schuck, I. Soysal, A. L. Todsen, M. R. Tünte, M. Vdovic, J. Vintr, M. Vovko, M. A. Vranka, L. Wagner, L. Wilkins, M. Willems, E. Wisdom, A. Yosifova, S. Zeng, M. A. Ahmed, T. Dwarkanath, M. Cikara, J. Lees, T. Folke, The general fault in our fault lines. *Nat. Hum. Behav.* **5**, 1369–1380 (2021).
111. C. A. Kelly, T. Sharot, Individual differences in information-seeking. *Nat. Commun.* **12**, 7062 (2021).
112. W. J. Brady, J. C. Jackson, B. Lindström, M. Crockett, *Preprint at OSF preprints* (2023); <https://doi.org/10.31219/osf.io/yw5ah>.
113. N. Ram, X. Yang, M. J. Cho, M. Brinberg, F. Muirhead, B. Reeves, T. N. Robinson, Screenomics: A new approach for observing and studying individuals' digital lives. *J. Adolesc. Res.* **35**, 16–50 (2020).
114. M. Brinberg, N. Ram, X. Yang, M. J. Cho, S. S. Sundar, T. N. Robinson, B. Reeves, The idiosyncrasies of everyday digital lives: Using the human screenome project to study user behavior on smartphones. *Comput. Hum. Behav.* **114**, 106570 (2021).
115. 2017 movement strategy. <https://meta.wikimedia.org/wiki/Strategy/Wikimedia-movement/2017/Direction>.
116. D. Dimitrov, F. Lemmerich, F. Flöck, M. Strohmaier, Different topic, different traffic: How search and navigation interplay on Wikipedia. *J. Web Sci.* **6**, 1–15 (2019).
117. I. Johnson, F. Lemmerich, D. Sáez-Trumper, R. West, M. Strohmaier, L. Zia, Global gender differences in wikipedia readership, in *Proceedings of the International AAAI Conference on Web and Social Media* (Association for the Advancement of Artificial Intelligence, 2021), vol. 15, pp. 254–265.
118. J. Mothe, G. Sahut, How trust in Wikipedia evolves: A survey of students aged 11 to 25. *Info. Res.* **23**, 783 (2018).
119. J. A. Litman, C. D. Spielberger, Measuring epistemic curiosity and its diverse and specific components. *J. Pers. Assess.* **80**, 75–86 (2003).
120. L. M. Giambra, C. J. Camp, A. Grodsky, Curiosity and stimulation seeking across the adult life span: Cross-sectional and 6- to 8-year longitudinal findings. *Psychol. Aging* **7**, 150 (1992).
121. H. I. Day, Hy I. Day, D. E. Berlyne, D. E. Hunt, Intrinsic motivation: A new direction in education (Holt, Rinehart and Winston of Canada, 1971).
122. P. H. Collins, *Intersectionality as Critical Social Theory* (Duke Univ. Press, 2019).
123. I. Kidd, J. M. James, G. P. Jr., Eds., *The Routledge Handbook of Epistemic Injustice* (Taylor and Francis, 2017).

124. A. J. van Deursen, E. Helsper, R. Eynon, J. A. van Dijk, The compoundness and sequentiality of digital inequality. *Int. J. Commun.* **11**, 452 (2017).
125. K. Manne, *Down Girl: The Logic of Misogyny* (Oxford Univ. Press, 2017).
126. I. Perry, *Vexy Thing* (Duke Univ. Press, 2018).
127. R. Su, J. Rounds, All STEM fields are not created equal: People and things interests explain gender disparities across STEM fields. *Front. Psychol.* **6**, 125967 (2015).
128. K. Weber, Gender differences in interest, perceived personal capacity, and participation in STEM-related activities. *J. Technol. Educ.* **24**, 18–33 (2012).
129. S. Kahn, D. Ginther, “Women and STEM,” Working paper, National Bureau of Economic Research, 2017.
130. T. B. Kashdan, R. A. Sherman, J. Yarbro, D. C. Funder, How are curious people viewed and how do they behave in social situations? From the perspectives of self, friends, parents, and unacquainted observers. *J. Pers.* **81**, 142–154 (2013).
131. P. H. Collins, Intersectionality and epistemic injustice, In *The Routledge handbook of epistemic injustice* (Routledge, 2017), pp. 115–124.
132. J. M. Box-Steffensmeier, J. Burgess, M. Corbetta, K. Crawford, E. Duflo, L. Fogarty, A. Gopnik, S. Hanafi, M. Herrero, Y. Y. Hong, Y. Kameyama, T. M. C. Lee, G. M. Leung, D. S. Nagin, A. C. Nobre, M. Nordentoft, A. Okbay, A. Perfors, L. M. Rival, C. R. Sugimoto, B. Tungodden, C. Wagner, The future of human behaviour research. *Nat. Hum. Behav.* **6**, 15–24 (2022).
133. A. Edelmann, T. Wolff, D. Montagne, C. A. Bail, Computational social science and sociology. *Annu. Rev. Sociol.* **46**, 61 (2020).
134. D. Lazer, E. Hargittai, D. Freelon, S. Gonzalez-Bailon, K. Munger, K. Ognyanova, J. Radford, Meaningful measures of human society in the twenty-first century. *Nature* **595**, 189–196 (2021).
135. D. M. J. Lazer, A. Pentland, D. J. Watts, S. Aral, S. Athey, N. Contractor, D. Freelon, S. Gonzalez-Bailon, G. King, H. Margetts, A. Nelson, M. J. Salganik, M. Strohmaier, A. Vespignani, C. Wagner, Computational social science: Obstacles and opportunities. *Science* **369**, 1060–1062 (2020).
136. H. Jach, R. Cools, A. Frisvold, M. Grubb, C. Hartley, J. Hartmann, L. Hunter, R. Jia, F. de Lange, R. Larisch, R. Lavelle-Hill, I. Levy, Y. Li, L. van Lieshout, K. Nussenbaum, S. Ravaioli, S. Wang, R. Wilson, M. Woodford, K. Murayama, J. Gottlieb, Individual differences in information demand have a low dimensional structure predicted by some curiosity personality traits. *Proc. Natl. Acad. Sci. U.S.A.* **1–64** (2024).
137. A. D. Sagar, A. Najam, The human development index: A critical review. *Ecol. Econ.* **25**, 249–264 (1998).
138. S. Alkire, Dimensions of human development. *World Dev.* **30**, 181 (2002).
139. United Nations Development Programme, *UNDP* (United Nations Development Programme, 2022).
140. P. Gildersleve, T. Yasseri, in *Complex Networks IX: Proceedings of the 9th Conference on Complex Networks CompleNet 2018 9* (Springer, 2018), pp. 271–282.
141. P. C. Terry, A. M. Lane, G. J. Fogarty, Construct validity of the profile of mood states-adolescents for use with adults. *Psychol. Sport Exerc.* **4**, 125–139 (2003).
142. Data platform/data lake/traffic/webrequest. <https://wikitech.wikimedia.org/wiki/Analytics/Data-Lake/Traffic/Webrequest>.
143. Data platform/data lake/traffic/botdetection. <https://wikitech.wikimedia.org/wiki/Data-Platform/Data-Lake/Traffic/BotDetection>.
144. X-analytics. <https://wikitech.wikimedia.org/wiki/X-Analytics>.
145. Wikimedia foundation country and territory protection list. <https://foundation.wikimedia.org/wiki/Legal:Country-and-Territory-Protection-List>.
146. A. Kline, Y. Luo, in *2022 44th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)* (IEEE, 2022), pp. 1354–1357.
147. A. Coutrot, E. Manley, S. Goodroe, C. Gahnstrom, G. Filomena, D. Yesiltepe, R. C. Dalton, J. M. Wiener, C. Hölscher, M. Hornberger, H. J. Spiers, Entropy of city street networks linked to future spatial navigation ability. *Nature* **604**, 104–110 (2022).
148. A. Coutrot, S. Schmidt, L. Coutrot, J. Pittman, L. Hong, J. M. Wiener, C. Hölscher, R. C. Dalton, M. Hornberger, H. J. Spiers, Virtual navigation tested on a mobile app is predictive of real-world wayfinding navigation performance. *PLOS ONE* **14**, e0213272 (2019).
149. G. Coughlan, A. Coutrot, M. Khondoker, A. M. Minihane, H. Spiers, M. Hornberger, Toward personalized cognitive diagnostics of at-genetic-risk Alzheimer’s disease. *Proc. Natl. Acad. Sci. U.S.A.* **116**, 9285–9292 (2019).
150. A. N. Rowan, K. Rowan, Highly pathogenic avian influenza. *WellBeing News* **5**, 1 (2023).
151. Manual:pagelinks table. <https://mediawiki.org/wiki/Manual:Pagelinks-table>.
152. A. Hagberg, P. Swart, D. S. Chult, “Exploring network structure, dynamics, and function using networkx” (Tech. Rep., Los Alamos National Lab, Los Alamos, NM, 2008).
153. M. Rubinov, O. Sporns, Complex network measures of brain connectivity: Uses and interpretations. *Neuroimage* **52**, 1059–1069 (2010).
154. T. P. Peixoto, The graph-tool python library (2014); <https://doi.org/10.6084/m9.figshare.1164194.v14>.
155. S. P. Borgatti, M. G. Everett, Models of core/periphery structures. *Soc. Net.* **21**, 375 (2000).
156. T. P. Peixoto, Hierarchical block structures and high-resolution model selection in large networks. *Phys. Rev. X* **4**, 011047 (2014).
157. P. D. Grünwald, *The Minimum Description Length Principle* (MIT Press, 2007).
158. A. Joulin, E. Grave, P. Bojanowski, M. Douze, H. J. Jégou, T. Mikolov, FastText.zip: Compressing text classification models. arXiv:1612.03651 [cs.CL] (2016).
159. G. Viswanathan, V. Afanasyev, S. V. Buldyrev, S. Havlin, M. G. E. da Luz, E. P. Raposo, H. E. Stanley, Lévy flights in random searches. *Phys. A Stat. Mech. Appl.* **282**, 1–12 (2000).
160. S. Maslov, K. Sneppen, Specificity and stability in topology of protein networks. *Science* **296**, 910–913 (2002).
161. P. Erdős, A. Rényi, On random graphs I. *Publ. Math. Debrecen* **6**, 18 (1959).
162. R. Albert, A.-L. Barabási, Statistical mechanics of complex networks. *Rev. Mod. Phys.* **74**, 47–97 (2002).
163. P. Révész, *Random Walk in Random and Non-Random Environments* (World Scientific, 2013).
164. R. Sinatra, J. Gómez-Gardenes, R. Lambiotte, V. Nicosia, V. Latora, Maximal-entropy random walks in complex networks with limited information. *Phys. Rev. E* **83**, 030103 (2011).
165. L. McInnes, J. Healy, J. Melville, UMAP: Uniform manifold approximation and projection for dimension reduction. arXiv:1802.03426 [stat.ML] (2018).
166. S. M. Mitchell, S. Lange, H. Brus, Gendered citation patterns in international relations journals. *Intl. Stud. Perspect.* **14**, 485–492 (2013).
167. M. L. Dion, J. L. Sumner, S. M. Mitchell, Gendered citation patterns across political science and social science methodology fields. *Polit. Anal.* **26**, 312–327 (2018).
168. N. Caplar, S. Tacchella, S. Birrer, Quantitative evaluation of gender bias in astronomical publications from citation counts. *Nat. Astron.* **1**, 0141 (2017).
169. D. Maliniak, R. Powers, B. F. Walter, The gender citation gap in international relations. *Intl. Org.* **67**, 889–922 (2013).
170. J. D. Dworkin, K. A. Linn, E. G. Teich, P. Zurn, R. T. Shinohara, D. S. Bassett, The extent and drivers of gender imbalance in neuroscience reference lists. *Nat. Neurosci.* **23**, 918–926 (2020).
171. M. A. Bertolero, J. D. Dworkin, S. U. David, C. L. Lloreda, P. Srivastava, J. Stiso, D. S. Bassett, Racial and ethnic imbalance in neuroscience reference lists and intersections with gender. bioRxiv. 336230 (2020). <https://doi.org/10.1101/2020.10.12.336230>.
172. X. Wang, J. D. Dworkin, D. Zhou, J. Stiso, E. B. Falk, D. S. Bassett, P. Zurn, D. M. Lydon-Staley, Gendered citation practices in the field of communication. *Ann. Int. Commun. Assoc.* **45**, 134–153 (2021).
173. P. Chatterjee, R. M. Werner, Gender disparity in citations in high-impact journal articles. *JAMA Netw. Open* **4**, e2114509 (2021).
174. J. M. Fulvio, I. Akinnola, B. R. Postle, Gender (Im)balance in citation practices in cognitive neuroscience. *J. Cogn. Neurosci.* **33**, 3–7 (2021).
175. D. Zhou, E. J. Cornblath, E. G. Teich, J. D. Dworkin, Gender diversity statement and code notebook v1.0 (2020).
176. A. Ambekar, C. Ward, J. Mohammed, S. Male, S. Skiena, in *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Association for Computing Machinery, 2009), pp. 49–58.
177. G. Sood, S. Laohaprapanon, Predicting race and ethnicity from the sequence of characters in a name. arXiv:1805.02109 [stat.AP] (2018).

Acknowledgments: We would like to thank L. Zia for insightful discussions and for reviewing an initial draft of this paper. **Funding:** 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.-S. acknowledges support from the National Institute on Drug Abuse (K01 DA047417). **Author contributions:** Writing—original draft: D.Z., D.M.L.-S., P.Z., M.G., and D.S.B. Conceptualization: D.Z., S.P., D.M.L.-S., P.Z., M.G., and D.S.B. Investigation: D.Z., S.P., D.M.L.-S., P.Z., M.G., and D.S.B. Writing—review and editing: D.Z., D.M.L.-S., M.G., P.Z., and D.S.B. Methodology: D.Z., S.P., D.M.L.-S., P.Z., M.G., and D.S.B. Resources: D.Z., S.P., D.M.L.-S., P.Z., M.G., and D.S.B. Data curation: D.Z., D.M.L.-S., and M.G. Validation: D.Z. and M.G. Formal analysis: D.Z., S.P., and M.G. Software: D.Z., S.P., and M.G. Project administration: D.Z., M.G., and D.S.B. Visualization: D.Z. **Competing interests:** The authors declare that they have no competing interests. **Data and materials availability:** All data needed to evaluate the conclusions in the paper are present in the paper, the Supplementary Materials, and in the permanent repository <https://zenodo.org/records/13922132>. The code used to process and aggregate the data to produce the results has also been deposited in the repository. Server logs may contain sensitive information with implications for the privacy of readers. Access to individual user data must be requested from the Wikimedia Foundation (https://foundation.wikimedia.org/wiki/Legal:Requests_for_user_information_procedures_and_guidelines). Pageview and clickstream data can be downloaded at <https://dumps.wikimedia.org/other/analytics/>.

Submitted 4 December 2023
Accepted 23 September 2024
Published 25 October 2024
10.1126/sciadv.adn3268