

Curious Rhythms: Temporal Regularities of Wikipedia Consumption

Tiziano Piccardi^{1*}, Martin Gerlach², Robert West^{3†}

¹Stanford University

²Wikimedia Foundation

³EPFL

piccardi@stanford.edu, mgerlach@wikimedia.org, robert.west@epfl.ch

Abstract

Wikipedia, in its role as the world’s largest encyclopedia, serves a broad range of information needs. Although previous studies have noted that Wikipedia users’ information needs vary throughout the day, there is to date no large-scale, quantitative study of the underlying dynamics. The present paper fills this gap by investigating temporal regularities in daily consumption patterns in a large-scale analysis of billions of timezone-corrected page requests mined from English Wikipedia’s server logs, with the goal of investigating how context and time relate to the kind of information consumed. First, we show that even after removing the global pattern of day–night alternation, the consumption habits of individual articles maintain strong diurnal regularities. Then, we characterize the prototypical shapes of consumption patterns, finding a particularly strong distinction between articles preferred during the evening/night and articles preferred during working hours. Finally, we investigate topical and contextual correlates of Wikipedia articles’ access rhythms, finding that article topic, reader country, and access device (mobile vs. desktop) are all important predictors of daily attention patterns. These findings shed new light on how humans seek information on the Web by focusing on Wikipedia as one of the largest open platforms for knowledge and learning, emphasizing Wikipedia’s role as a rich knowledge base that fulfills information needs spread throughout the day, with implications for understanding information seeking across the globe and for designing appropriate information systems.

Introduction

Human life is driven by strong temporal regularities at multiple scales, with events and activities recurring in daily, weekly, monthly, yearly, or even longer periods (Kreitzman and Foster 2011). The cyclical nature of human life is also reflected in human behavior on digital platforms. Understanding and modeling temporal regularities in digital-platform usage is important from a practical as well as a scientific perspective: on the practical side, understanding user behaviors and needs is critical for building effective, user-friendly platforms; on the scientific side, since online behavior reflects user needs, studying temporal regularities of

platform usage can shed light on the structure of human life itself and is thus consequential for sociology, psychology, anthropology, medicine, and other disciplines. For instance, in information science, time has long been recognized as a crucial contextual factor that drives human information seeking (Savolainen 2006). The study of traces recorded on digital platforms has yielded novel insights about circadian rhythms (Aledavood et al. 2022; Pintér and Felde 2022; Althoff et al. 2017; Murnane et al. 2015; Chalmers et al. 2011; Kumar Swain and Kumar Pati 2020; Leypunskiy et al. 2018) and periodic fluctuations in alertness (Abdullah et al. 2016), mood (Dodds et al. 2011; Golder and Macy 2011), focus (Mark et al. 2016), musical taste (Park et al. 2019; Heggli, Stupacher, and Vuust 2021), purchase habits (Gullo et al. 2019), and ad engagement (Saha et al. 2021).

With regard to both aspects—practical and scientific—Wikipedia constitutes a particularly important online platform to study: On the practical side, Wikipedia is one of the most frequently visited sites on the Web, such that a better understanding of user behavior can potentially lead to site improvements with consequences for billions of users. On the scientific side, Wikipedia is not just yet another website; it is the world’s largest encyclopedia, where each page is about a precise concept. Temporal regularities of Wikipedia usage thus have the potential to reveal regularities of human necessities, telling us what humans care about, and what they are curious about, at what times. Some temporal regularities of Wikipedia usage are already known: e.g., Wikipedia’s overall usage frequency, as well as the length of sessions, varies by time of day (Piccardi et al. 2023), as do users’ reasons for reading Wikipedia (Singer et al. 2017; Lemmerich et al. 2019). These studies have, however, glanced at temporal regularities merely superficially and in passing while focusing primarily on other aspects of Wikipedia usage. The most related to our work is a study that showed that Wikipedia *editing* follows daily and weekly regularities (Yasseri et al. 2012), but the study was limited to editing behavior. One reason why previous studies have not been able to analyze reading, rather than editing, behavior to date is that, as opposed to edit logs (where editors’ geo-locations can be approximated via logged IP addresses), geo-located reading logs are not publicly avail-

*Work done while at EPFL.

†Robert West is a Wikimedia Foundation Research Fellow.
Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

able. Rather, Wikipedia’s public hourly pageview logs¹ report only Coordinated Universal Time (UTC) without specifying the reader’s local time. Especially in large language editions such as English, which are accessed from many countries in different timezones, this constitutes a major limitation.

We overcome this limitation by working with a non-public dataset of English Wikipedia’s full webrequest logs, enriched with timezone information inferred from user IP addresses, which allows us to timezone-correct all timestamps and thus to faithfully study, for the first time, temporal regularities of Wikipedia reading at hourly granularity. We characterize how information consumption on the platform varies by time of day and how time interacts with other contextual properties, such as article topics and reader country. Adding to previous studies on the consumption and popularity of Wikipedia content, we provide insights into the daily temporal rhythms that drive Wikipedia usage. Specifically, we address the following research questions:

RQ1 Strength of rhythms: How strongly is Wikipedia consumption driven by periodic rhythms?

RQ2 Shapes of rhythms: What are the typical shapes of Wikipedia consumption rhythms?

RQ3 Correlates of rhythms: How do topical and contextual factors determine Wikipedia consumption rhythms?

Regarding RQ1, we find that fluctuations in Wikipedia’s total access volume are largely explained by a diurnal baseline rhythm corresponding to the human circadian wake–sleep cycle. Crucially, however, individual articles deviate systematically from the baseline rhythm, with deviations themselves following periodic diurnal patterns.

Regarding RQ2, principal component analysis reveals that individual articles’ access rhythms are heavily driven by a small number of prototypical temporal signatures, but that article-specific rhythms do not cluster into distinct groups, but vary smoothly along a continuum.

Regarding RQ3, regression analysis shows that temporal regularities in article access volume vary systematically by article topic, access method (mobile vs. desktop), and reader country. The latter is the strongest determinant of access rhythm, whereby different countries’ interests in English Wikipedia fluctuate in idiosyncratic periodic patterns over the course of the day. In terms of topic, to a first approximation, STEM and HISTORY & SOCIETY articles tend to be more popular early in the day, and MEDIA articles later in the day, with CULTURE articles being less concentrated in time.

Taken together, these findings further our understanding of how human information needs vary in space and time, and can serve as a stepping stone toward the informed design of improved information systems, on Wikipedia and beyond.

Related Work

Temporal rhythms in digital traces. Temporal patterns using digital traces have been explored in a large variety

of topics. Consumption logs can bring invaluable insight into our understanding of how time impacts human activities. For example, mobile phone traces can help characterize daily rhythms (Aledavood et al. 2022) and expose insights about our circadian rhythm (Pintér and Felde 2022), which may be difficult to study otherwise. Analyzing the temporal dynamics of digital traces can also uncover biological and cognitive rhythms during the day, such as sense of alertness (Abdullah et al. 2016), chronotypes (Althoff et al. 2017), wellness (Zhou et al. 2023b), and focus (Mark et al. 2016). Individual temporal rhythms can also serve as predictors of medical conditions that need attention (Doryab et al. 2019). Social media data, like Twitter (Murnane et al. 2015; Chalmers et al. 2011), and messaging apps, like WhatsApp (Kumar Swain and Kumar Pati 2020), can be used as large-scale sensors revealing insights about sleeping patterns (Leypunskiy et al. 2018), and happiness fluctuations (Dodds et al. 2011), with variations across cultures (Golder and Macy 2011). Time of day also impacts our choices in terms of musical taste (Park et al. 2019; Heggli, Stupacher, and Vuust 2021), purchase habits (Gullo et al. 2019), and ad engagement (Saha et al. 2021).

Interactions with Wikipedia, too, are affected by time. Previous studies describe how consumption patterns reveal information about seasonal fluctuations in mood (Dzogang, Lansdall-Welfare, and Cristianini 2016) at the population scale and how editors’ temporal patterns exhibit culture-specific differences (Yasseri, Sumi, and Kertész 2012). It has also been noted that modeling Wikipedia temporal readers’ needs has implications for technical improvements and could be exploited to develop scalable infrastructure adapted to reading patterns (Reinoso et al. 2012).

Wikipedia reader behavior. Given Wikipedia’s central role in the Web ecosystem, there is increasing attention to characterizing reader behavior. Previous work investigated the reasons that lead readers to consume Wikipedia content (Singer et al. 2017), finding differences by time and country (Lemmerich et al. 2019). Complementary studies focus on consumption dynamics (Ratkiewicz et al. 2010; Ribeiro et al. 2021; Lehmann et al. 2014; TeBlunthuis, Bayer, and Vasileva 2019) and interaction with various article elements, such as citations (Maggio et al. 2020; Piccardi et al. 2020), external links (Piccardi et al. 2021), and images (Rama et al. 2021). Significant attention is also dedicated to navigating from page to page to understand the mechanism that allows users to move across content in a natural setup (Piccardi et al. 2023; Piccardi, Gerlach, and West 2022; Trattner et al. 2012; Helic et al. 2013; Singer et al. 2014; Lamprecht et al. 2017; Gildersleve and Yasseri 2018; Dimitrov et al. 2019) or in lab-based settings, such as wiki games (West, Pineau, and Precup 2009; West and Leskovec 2012b; Helic 2012; West and Leskovec 2012a; Scaria et al. 2014; Koopmann et al. 2019). Our work complements these findings by focusing on the shape of temporal regularities, which have not been analyzed in detail before.

Information needs and curiosity. Our work relates broadly to more theoretical work on information needs. The formulation of information needs was initially popularised by Wil-

¹<https://dumps.wikimedia.org/other/pageviews/readme.html>

son (Wilson 1981) and the resulting information models refined over the years (Wilson 1997, 1999). Information needs have also been linked to concepts in biology by comparing the need for information to the need for food (Machlup 1983). This idea is formalized in information foraging theory (Pirolli and Card 1999), which has developed behavioral models describing humans in the information space as predators relying on mechanisms such as information scent (Chi et al. 2001) to find what they need (Suh et al. 2010; Mangel et al. 2013; Rotman et al. 2011). More recent work (Lydon-Staley et al. 2021; Zhou et al. 2023a) investigates the mechanism of online consumption by focusing on the role of curiosity, using Wikipedia as the reference platform and finding substantial differences in how humans explore information networks.

Data

Our study relies on the access logs of English Wikipedia, collected over four weeks (1–28 March 2021) on the servers of the Wikimedia Foundation. These server logs, stored for a limited time, describe the requests received by the server when readers access the site, capturing the title of the requested article, the time the request was made, the user’s IP address and geo-location (approximately inferred from the IP address), and more.

In this study, we focus only on the traffic generated by human readers by discarding all bots. In order to ensure anonymity, we prepare the data as follows. First, we consider only requests for articles (namespace 0), ignoring requests for pages in other namespaces (e.g., talk pages), and we consider only requests originating from external websites, ignoring requests made from other Wikipedia pages. The restriction to incoming traffic from external websites better captures the (exogenous) information needs that cause users to visit Wikipedia, rather than (endogenous) information needs that are caused by visiting Wikipedia. Additionally, we refine the logs by removing sequential loads of the same page from the same client because they could be an artifact of the browser (Piccardi et al. 2023), not representing a real intention to visit Wikipedia.

We anonymize the data by removing all sensitive information such as IP addresses, user-agent strings, fine-grained geo-coordinates, and all requests from logged-in users (3% of the pageloads). Finally—and crucially for our purposes—we align all requests by converting timestamps to the user’s local time using timezone information available in the logs.

After these steps, we retain 3.45B pageloads associated with 6.3M articles. We represent the number of pageloads of article a in hour h of the week (averaged over the four weeks) by $n_a(h)$, i.e., each article is represented by a 168-dimensional time series, with one entry for each of the $168 = 24 \times 7$ hours per week.

Article properties. ORES (Halfaker and Geiger 2019), Wikipedia’s official article scoring system, offers a classifier for labeling articles with topics. The labels are organized in a two-level taxonomy manually derived from WikiPro-




Figure 1: *Top*: Total number of pageloads from external origin by hour of the week, averaged over four weeks. *Bottom*: Idem, stratified by desktop and mobile.

jects² (groups of editors who self-organize to work on specific topical areas). We applied the classifier to all articles in our dataset and, for each article, obtained a probability for each of 38 topics, grouped into five top-level topics: STEM, CULTURE, HISTORY & SOCIETY, MEDIA (originally included in CULTURE, but promoted to a top-level topic for our study), GEOGRAPHICAL. Note that each article is independently scored for each topic, so topic probabilities do not sum to 1 for a given article.

Access properties. For each pageload, we retain two contextual properties of the request: access method and user country. The access method indicator specifies if the user loaded the article from a desktop or a mobile device, such as a smartphone or tablet. The user country is estimated by geo-locating the IP address associated with the request.

Baseline rhythm. Fig. 1 (top) shows the total number of pageloads per hour of the week (averaged over the four weeks considered). The consumption rhythm follows a diurnal pattern, with the lowest access volume between 4:00 and 5:00 from Monday to Saturday and between 5:00 and 6:00 on Sunday, and with a sharp increase in traffic from 18:00 every day of the week, consistently reaching its peak at 21:00. When broken down by access method (Fig. 1, bottom), different device types are associated with different patterns, with mobile devices driving the increase in evening traffic. Similarly, access from desktop devices, likely more dependent on working rhythms, shows a small reduction around 12:00 and 18:00 and reduced activity on weekend mornings.

Denoting the number of pageloads in hour h by $N(h) = \sum_a n_a(h)$ and normalizing $N(h)$ to be a distribution over the 168 hours of a week, we obtain what we term Wikipedia’s *baseline rhythm* $\Pr(h) := N(h) / \sum_{h'=0}^{167} N(h')$.

Divergence from the baseline rhythm. In addition to Wikipedia’s overall access volume (Fig. 1), we are interested in studying temporal access volume regularities for individual articles a , by analyzing article-specific distributions over hours, $\Pr(h|a) := n_a(h) / \sum_{h'=0}^{167} n_a(h')$. Individual articles’ rhythms $\Pr(h|a)$ are heavily driven by the overall base-

²<https://en.wikipedia.org/wiki/Wikipedia:WikiProject>

line rhythm $\Pr(h)$ dictated by human wake–sleep patterns. Therefore, to study article-specific rhythms in isolation from the baseline rhythm, we remove the latter by computing the hourly *divergence* $D_a(h)$ of a 's rhythm $\Pr(h|a)$ from the baseline rhythm $\Pr(h)$ via pointwise division:

$$D_a(h) = \frac{\Pr(h|a)}{\Pr(h)} = \frac{n_a(h) / \sum_{h'=0}^{167} n_a(h')}{N(h) / \sum_{h'=0}^{167} N(h')}. \quad (1)$$

A divergence greater [less] than 1 indicates that in hour h , article a receives more [less] attention than expected from the baseline rhythm, whereas a divergence of 1 implies that a global weekly rhythm fully determines a 's consumption.

Given the skewed distribution of article popularity, to avoid data sparsity issues, we keep only articles with at least 1,000 pageloads over the four weeks considered, or 35 pageloads per day on average. After applying this filter, we retain 439K articles, accounting for 83.8% (2.89B pageloads) of English Wikipedia's traffic.

RQ1: Strength of Rhythms

We start by quantifying the strength of temporal regularities in Wikipedia consumption patterns. Taking a signal processing approach, we decompose a time series $x(h)$ of interest (e.g., $x(h) = \Pr(h)$ or $D_a(h)$) into its frequency components via the Fourier transform, which represents $x(h)$ as a weighted sum of sinusoids of all possible frequencies. We denote the weight (Fourier coefficient) of frequency $f \in \{0, \dots, 167\}$ by $\hat{x}(f)$. We then measure the contribution of each frequency f to time series $x(h)$ (or, equivalently, to $\hat{x}(f)$) via the so-called *energy spectral density* $E(f)$, which captures the fraction of $x(h)$'s total variance (or energy) explained by each frequency f :

$$E(f) = \frac{|\hat{x}(f)|^2}{\sum_{f'=0}^{167} |\hat{x}(f')|^2}. \quad (2)$$

Baseline rhythm. Using $x(h) = \Pr(h)$ lets us study temporal regularities in the weekly baseline rhythm. Fig. 2 (top) plots the variance $E(f)$ explained by each frequency component f of the baseline rhythm $\Pr(h)$ (mean-centered across h in order to discard the constant offset term corresponding to $f = 0$). As previously anticipated by Fig. 1, the dominant frequency corresponds to a 24-hour cycle ($f = 7$). The daily and half-daily cycles ($f = 7$ and 14; wavelengths 24 and 12 hours) together explain 96.2% of the variance (74.8% and 21.4%, respectively). When split by access method, the pattern on mobile is more predictable, with daily and half-daily cycles explaining 94% of the total variance for mobile, vs. 87% for desktop.

Article-specific rhythms. To measure the strength of article-specific daily rhythms that do not depend on the overall baseline rhythm $\Pr(h)$, we next repeat the above analysis, but now with $x(h) = D_a(h)$ (again mean-centered across h). We obtain the energy spectral density $E(f)$ for each article a and average it over a in order to measure the average strength of frequency f in divergences $D_a(h)$ from the baseline rhythm, plotted in Fig. 2 (bottom). We see that, even after removing




Figure 2: *Top/blue*: Contribution of each frequency to baseline rhythm $\Pr(h)$ of Wikipedia access volume (measured as the fraction of total variance explained). *Bottom/red*: Contribution of each frequency to article-specific divergence $D_a(h)$ from baseline rhythm (Eq. 1) (computed per article a , then averaged over articles).

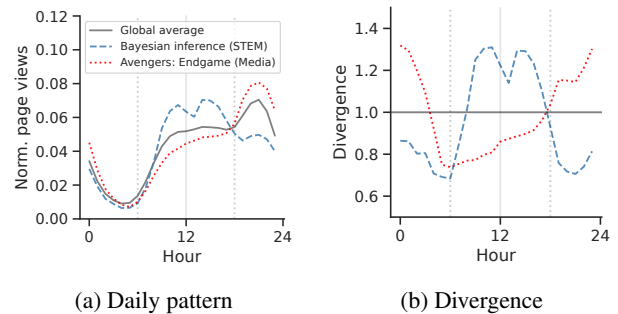


Figure 3: Daily access volume of two articles with different topics: STEM (dashed blue) and MEDIA (dotted red). (a) Normalized time series $\Pr(h|a)$. (b) Divergence $D_a(h)$ of $\Pr(h|a)$ from the overall baseline rhythm $\Pr(h)$ (cf. Eq. 1).

the overall baseline rhythm, still 22.3% of the variance is explained by a combination of the cycles of wavelengths of 24 (16.5%) and 12 (5.8%) hours. This means that the individual articles' access volume rhythms are not just driven by the change in the overall number of pageloads throughout the day/week, but are also strongly determined by article-specific factors.

To summarize, investigating RQ1, we found that overall Wikipedia consumption follows a strong baseline day–night rhythm, but also that individual articles considerably deviate from this baseline rhythm. The deviations themselves also largely follow diurnal patterns.

RQ2: Shapes of Rhythms

Answering RQ1, we established that the consumption pattern of articles exhibits daily regularities even after removing the weekly baseline rhythm. Next, we investigate the shape of these patterns. Since the daily rhythm (wavelength 24 hours) was found to be the strongest by far, we henceforth focus on the daily instead of the weekly cycle; i.e., we now have $h \in \{0, \dots, 23\}$ instead of $\{0, \dots, 167\}$, with averages taken over 28 days instead of over four weeks.




Figure 4: Four principal components of the daily access volume time series, capturing 73.6% of total variance.

Fig. 3 shows examples of the resulting daily time series for two articles a associated with different topics (STEM and MEDIA). Fig. 3a shows the shape of the normalized daily access volume time series $\Pr(h|a)$, and Fig. 3b shows their divergence $D_a(h)$ from the daily baseline rhythm $\Pr(h)$.

Principal Components

To investigate prototypical article consumption behaviors, we extract the principal components describing the daily divergence time series $D_a(h)$ of article a . For this purpose, we stack the individual divergence time series $D_a(h)$ in a $439K \times 24$ matrix D whose rows correspond to a and whose columns correspond to h . We mean-center D column-wise and compute its singular value decomposition (SVD) $D = U\Sigma V^T$. The (orthogonal) columns of V , then, are the principal components of D , capturing prototypical divergence time series. Individual articles' divergence time series $D_a(h)$ can be approximated as linear combinations of the top principal components. Σ is a diagonal matrix whose i -th entry represents the standard deviation of the data points (rows of D) when projected on the i -th principal component. Using the elbow method, we find that the first four principal components provide a good representation of consumption behavior, accounting, respectively, for 39.2%, 19.7%, 10.2%, and 4.5% of the total variance, jointly capturing 73.6% of the total variance.

Fig. 4 shows the first four principal components of D . Since each component can contribute positively or negatively to the reconstruction of the rows of D (since principal components are sign-invariant), we depict each component in its positive and negative variant (dark and light). The first and strongest component (PC1) captures the distinction between articles that receive more attention during the day versus evening/night (or vice versa), with two turning points around 7:00 and 21:00. The second component (PC2) describes patterns with a strong positive or negative contribution in the morning, with a peak around 5:00, which is reversed in the late hours of the day. The third principal component (PC3) captures a consumption pattern that peaks positively or negatively in the early morning and evening. Finally, the fourth principal component (PC4) models a similar behavior with the strongest contribution around 4:00 and in the early evening.

Clustering

Next, we aim to identify groups of articles with similar temporal access patterns via clustering. As clustering tends to work better in low dimensions, we start by obtaining low-dimensional time series representations by truncating U (obtained via SVD) to the first four columns $U_{1:4}$ (corresponding to projections on the first four principal components) and approximating article a 's divergence time series $D_a(h)$ via the a -th row of $U_{1:4}$. We then cluster the rows of $U_{1:4}$ using k -means and search for the optimal number k of clusters using the average silhouette width and sum-of-squares criteria. Both criteria indicate $k = 2$ as the optimal number of clusters, suggesting that the articles cannot easily be separated into distinct groups. This intuition is supported by Fig. 5a, which shows the clusters obtained for $k = 2, 3, 4$. The upper plots show a UMAP (McInnes, Healy, and Melville 2018) reduction in two dimensions of the four principal components used to cluster the articles. The visualization is based on a sample of 10K random articles, color represents cluster assignments, and centroids are marked as black dots. The lower plots show the centroid of each cluster, reconstructed from the first four principal components. These plots support the intuition that consumption behaviors do not separate into different groups but distribute along a continuum.³

On the right side of each of the three scatter plots, we see articles popular during daytime. In contrast, on the left, we see articles popular during the evening/night. This partition becomes more evident when increasing the number of clusters. With three clusters, the data on the left separate into an additional group (green) containing articles popular during the night. Similarly, adding the fourth cluster isolates a set of articles (red) popular in the early morning.

To summarize, investigating RQ2, we found that Wikipedia articles' access rhythms are driven by only few prototypical basis patterns (e.g., day vs. night). Individual articles' access patterns can be captured as weighted combinations of those basis patterns, yet they do not cluster into distinct, well-separated groups, but vary along a smooth continuum.

RQ3: Correlates of Rhythms

In this section, we investigate what factors correlate with the shape of daily consumption. We focus on three factors: article topics, access method, and user country.

Fig. 5b offers initial evidence that article topics are associated with daily access patterns. The plot represents the subset of articles from Fig. 5a with topic MEDIA or STEM. By coloring the data by topic (red for MEDIA, blue for STEM), we can notice a natural separation akin to the clustering with $k = 2$ clusters (Fig. 5a, left). Note that the topic was not considered during clustering; the separation happens only based on the shape of access patterns. Articles about STEM are on the right side, indicating a pattern aligned with more attention during daytime, whereas articles about MEDIA cover the right side, associated with evening and night consumption.

³Silhouette scores obtained by grid-searching the number of principal components and clusters do not improve the separation. Density-based clustering yields the same conclusions.

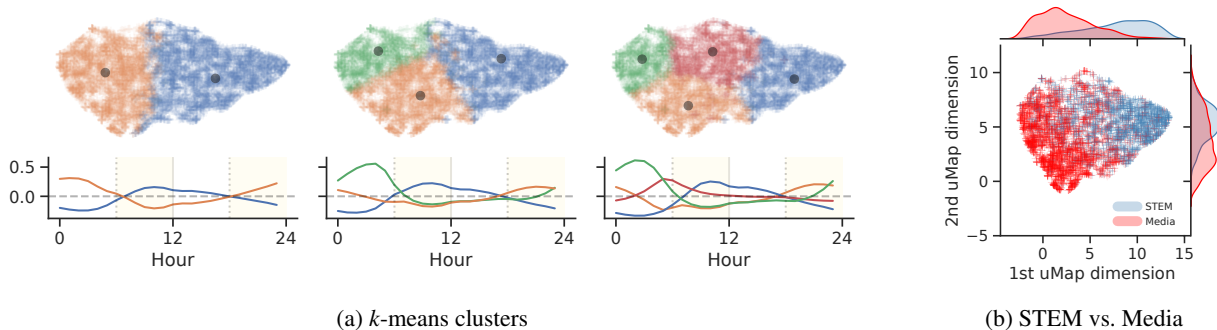


Figure 5: (a) *Top*: UMAP projection of access volume time series for 10K random articles, with colors representing clusters obtained with k -means for different values of $k = 2, 3, 4$. *Bottom*: Centroids reconstructed from their first four principal components. (b) A subset of the same 10K articles, with colors representing topics (STEM and MEDIA).

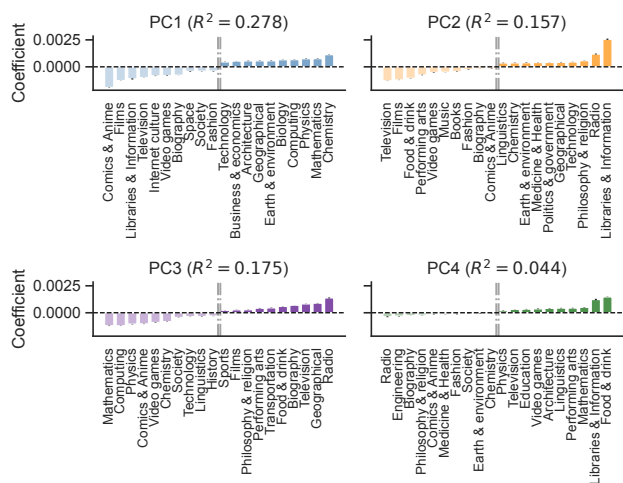


Figure 6: Linear regression coefficients of the topics most associated with each of the four main principal components of article access volume time series.

Going deeper, we analyze prototypical shapes of access volume time series via regression analysis, quantifying the relationship of each factor (topics, access method, country) with time, and measuring the influence of each factor in an ablation study.

Principal Components and Topics

In an initial analysis, we leverage the principal components introduced earlier for RQ2 and investigate the relationship between topics and access patterns using four linear regressions, one per top principal component. Given an article a , we predict the projection U_{ai} of its divergence time series $D_a(h)$ onto the i -th principal component using a 's topics as binary predictors.

Fig. 6 shows a summary of the coefficients obtained by fitting the four regressions. The first principal component, capturing the day–night rhythm, separates articles about STEM from those about entertainment (also cf. Fig. 5b). Articles associated with topics such as COMICS & ANIME, FILMS, and TELEVISION show a negative association with this prin-

cipal component, suggesting nighttime popularity, whereas the temporal patterns of articles about CHEMISTRY, MATHEMATICS, and PHYSICS show a positive association, suggesting daytime popularity. The second component, associated with higher activity in the early morning and a drop during the evening, shows that the topics associated most positively with this pattern are LIBRARIES & INFORMATION, RADIO, and PHILOSOPHY & RELIGION. On the other side of the spectrum, FILMS, TELEVISION, and FOOD & DRINKS are most negatively associated with this pattern (i.e., low activity in the early morning, spike during the evening). Finally, the third and fourth principal components, showing peaks in the early morning and during the evening, are most positively associated with RADIO and FOOD & DRINK, respectively.

Temporal Variation of Correlates

Next, we are interested in the prototypical access volume time series associated with each level of each factor (topics, access method, user country). A naive way of accomplishing this would be to first restrict the data to the respective factor level (e.g., requests for BIOLOGY articles only, or requests from Canada only) and then compute access volume time series for the restricted data only. This approach is, however, compromised by the fact that the various factors contribute differently to the observed access pattern. Thus, in order to disentangle the factors, we employ linear regression.

Model. We prepare the data as follows. First, we decompose each article a 's divergence time series $D_a(h)$ by country and access method. To avoid data sparsity issues, we limit ourselves to the 20 countries with the most pageloads. This step generates, for each article, 40 daily time series (20 countries, two access methods) describing the divergence of the attention to a from a specific country and access method vs. the global baseline rhythm $\Pr(h)$. We explode each time series into 24 samples, such that each article is represented by 960 data points describing its divergence from the baseline rhythm for each combination of 20 countries, two access methods, and 24 hours. We then model divergence (transformed via the natural logarithm, to make the model multiplicative) as a linear combination of country, access method, topics, and hour of the day. Additionally, since we are in-




Figure 7: Linear regression coefficients of the interaction between topics and hour. Background colors represent top-level topics: STEM (green), MEDIA (yellow), CULTURE (blue), HISTORY & SOCIETY (red) and GEOGRAPHICAL (gray). For each top-level topic, specific topics are sorted alphabetically. Red bands represent 95% confidence intervals.




Figure 8: Linear regression coefficients of the interaction between country and hour (left, gray background, sorted by the total access volume), and between device and hour (rightmost plot). Yellow area in rightmost plot represents typical working hours in Western countries. Red bands represent 95% confidence intervals.

terested in modeling the relationship between country, access method, and topics *with time*, we include interaction terms between each of these three factors with the hour-of-the-day indicator. We represent all binary predictors via deviation coding, which lets us interpret the coefficients associated with each level as differences from the grand mean, which is itself captured by the intercept (since differences are taken in log space, they correspond to ratios with the grand mean in linear space). With this setup, we keep the time series composed of at least 100 pageloads during the four-week period of study and fit a linear regression using a random sample of 30K articles. The obtained model fits the data with $R^2 = 0.181$.

Coefficient analysis. Sorting (by hour) the 24 interaction coefficients of each factor level with hour of the day allows us to characterize the typical temporal shape of each factor level (in terms of its deviation from Wikipedia’s baseline rhythm) *when controlling for all other factors*.

Fig. 7 shows the temporal shapes of the topics organized in the five top-level topical groups. The plot shows how articles about STEM topics, such as CHEMISTRY, PHYSICS, and MATHEMATICS, tend to receive more attention than average during the daytime and a visible reduction outside the typical working hours. On the other hand, articles about FILMS, TELEVISION, and BIOGRAPHY have an inverted shape, with less consumption than average during the day and a substantial increase during the evening. Interestingly, the shapes of the temporal patterns suggest that content about VIDEO GAMES, COMICS & ANIME, INTERNET CULTURE, MILITARY, and SOCIETY are consumed by night-active readers, with relative consumption peaking during the night.

On the contrary, articles about RADIO, LIBRARIES & INFORMATION, and PHILOSOPHY, are more preferred in the early hours of the day. Some of the shapes, especially the

ones associated with STEM articles, show a reduction of attention around noon, suggesting that they might be affected by the lunch break when people’s attention moves to other content types. This is corroborated by the fact that attention on articles about FOOD sees an increase during common meal times.

Next, Fig. 8 shows the interaction coefficients of country by time. Some countries share similar daily patterns. For example, readers from the United States, Germany, Netherlands, and Nigeria tend to consume Wikipedia more than average in the early morning. This behavior is inverted for readers from India, Ireland, Italy, and Spain, who during the same hours consume less content than average. Meanwhile, other countries, such as Malaysia, Singapore, Brazil, Russia, and Pakistan, show higher consumption during the night. Furthermore, some countries, like the Philippines, Italy, France, and Spain, also reveal shared habits, such as a reduction of information consumption around noon, possibly associated with lunchtime.

Finally, the rightmost plot of Fig. 8 shows the coefficients of the interaction of access method with time, in particular, the shape of the consumption from desktop devices. As already observed by breaking down the baseline rhythm by device in Fig. 1, access from desktop devices is above the global average in the central hours of the day, between 9:00 and 17:00.

Typical Access Times of Topics

We now aim to summarize each topic-wise time series of Fig. 7 concisely in a single number describing at what time of day articles from the respective topic are particularly popular. In order to obtain point estimates of “average time”, the simple arithmetic mean is unsuitable, given the circular nature of time (e.g., the average between 23:00 and 1:00

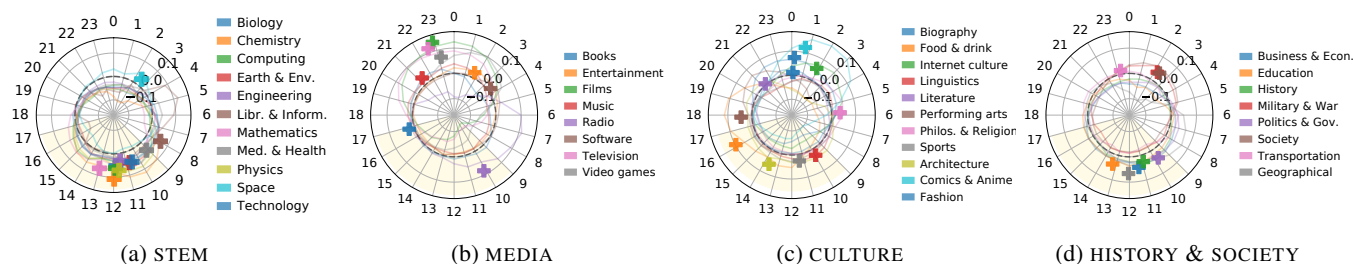


Figure 9: Typical access times of topics. Lines represent the coefficient time series of Fig. 7, and bold crosses the average times per topic. The dotted line marks the global baseline rhythm $Pr(h)$. Yellow area shows typical working hours in Western countries. GEOGRAPHICAL is included in HISTORY & SOCIETY to compress the visualization.

should be 0:00, not the arithmetic mean of 12:00), so we use the angular mean instead (essentially an arithmetic mean in the 2D plane described by a clock’s hands). Fig. 9 shows the topic-wise time series of Fig. 7 in a circular fashion as light-colored curves, and average times as bold crosses, with one panel per top-level topic (where a cross’s closeness to the origin captures the variance of the respective time series). We observe that articles about STEM (except for SPACE) tend to garner particularly much attention during working hours between 8:00 and 14:00. On the other hand, MEDIA is, on average, consumed more during the evening and night. RADIO is an exception in this group, peaking in the early morning, as visible in Fig. 7, bringing its daily average to around 10:00. Articles about BOOKS are at the limit of the typical working hours, with average access in the late afternoon around 17:00. Differently from the previous two groups, articles about CULTURE tend to have high variance and have average times spread over the day. Even in this case, entertainment topics such as FASHION, INTERNET CULTURE, and COMICS & ANIME are concentrated during the night, with average access after midnight. Finally, content about HISTORY & SOCIETY shows an average consumption split into two groups, night vs. working hours, with SOCIETY, MILITARY, and TRANSPORTATION concentrated during the late evening and night, and BUSINESS & ECONOMICS, HISTORY, and EDUCATION during the day.

Strength of Correlates: Ablation Study

Finally, we investigate each factor’s influence on temporal access rhythms by estimating the strength of the interaction between each factor and time in an ablation study. Using the fitted model described above, we proceed as follows. We use a held-out dataset of 10K random articles not used in model-fitting to assess the R^2 fit after permuting all the interaction terms of the factor being investigated. This approach, typically used to estimate feature importance in machine learning (Breiman 2001; Fisher, Rudin, and Dominici 2019), has the advantage of keeping the model fixed—thus, no R^2 adjustment is required—and measuring the impact of removing the correlation between the selected feature and the dependent variable.

The original model has an R^2 fit of 0.181. When permuting the interaction terms of “time by topic”, the R^2 drops to 0.055 (reduction by 69%); when permuting “time by access method”, the R^2 is 0.040 (reduction by 78%); and fi-

nally, when permuting “time by country”, the R^2 score is -0.0298^4 (reduction of 116%). A permutation of all three factors reduces the R^2 to -0.429 . These observations suggest that all factors are essential for the prediction and indicate that the interaction of time with the readers’ country plays the most important role.

To summarize, in answering RQ3, we found that temporal regularities in the popularity of Wikipedia articles vary in important and systematic ways by topic, access method, and user country, and that, out of these, user country has the strongest influence.

Discussion

Summary of findings. In this study, we conducted the first large-scale study on temporal rhythms of information consumption on Wikipedia, based on a new timezone-corrected dataset of hourly-aggregated pageviews where the reader’s local time was inferred from Wikipedia’s webrequest logs.

First, we showed that the overall information consumption exhibits daily rhythmic patterns, with the strongest components having periods of 24 and 12 hours, respectively. We further showed that the consumption of individual articles follows specific rhythms that cannot be explained by Wikipedia’s overall wake–sleep baseline rhythm. Rather, each article reveals a specific consumption fingerprint throughout the day.

Second, we provided a systematic description of the principal shapes of the different rhythms of the individual articles. We find that the main shape distinguishes articles that are read disproportionately more during the day than during the night (and vice versa). We do not, however, find distinct clusters of shapes but, instead, observe a continuum of different rhythms of information consumption.

Third, we show systematic differences in consumption patterns based on the reader’s context (country or device) or the article’s topics. We showed that articles of specific topics are associated with specific rhythmic patterns. For example, articles on STEM and MEDIA naturally separate into two clusters related to rhythms with disproportional attention during the day and night, respectively. More generally, this leads to markedly different “average” times

⁴A negative R^2 is possible on a held-out set, indicating that a flat line approximates the data better than the model.

when articles from different topics are accessed. We also showed that pageloads from mobile vs. desktop devices are driven by substantially different 24-hour rhythms, with mobile pageloads showing an almost twofold increase in the evening hours compared to the day, which is absent for desktop devices. Lastly, we also find substantial variation in the access patterns across different countries.

Implications. Our work shows that context is an important element to consider when trying to understand how information on Wikipedia is consumed. This has several implications.

Diversity of information needs. Wikipedia as a platform fulfills multiple information needs. These needs vary not only with geographic location, such as country (Lemmerich et al. 2019), but also with time of day when a reader accesses the platform. In order to serve these needs, we need to consider the heterogeneity of Wikipedia’s audience in *space and time*. Additionally, given the extension of topics that Wikipedia covers and its large adoption, our study offers insights into what content people consume online during the day, at a scale generally accessible only to search engines and Internet providers, with implications for design beyond Wikipedia.

Cultural diversity. Given the diversity of Wikipedia access rhythms across countries (Fig. 8), future work could revisit our results from an anthropological perspective, e.g., in order to pinpoint cultural differences in the daily rhythm of information needs—or curiosity—across the globe. For instance, we are fascinated by the fact that the rhythm of (English) Wikipedia’s popularity in the U.S. is nearly exactly opposed to that in Spain, and that India and Pakistan—two countries that once were one—have entirely different rhythms of English Wikipedia consumption. Wikipedia logs offer a window into where in the world people care about what, when, and we anticipate that studying access time series while accounting for the interaction of country by topic (which we omitted here) has enormous potential for understanding global patterns of needs for knowledge.

Metrics for information needs. The popularity of an article (i.e., pageloads) is often used as an indicator for its relevance (Wulczyn et al. 2016) or a covariate used for stratification in the analysis of articles (Piccardi et al. 2020). In our context, this corresponds to only looking at the volume of the consumption pattern and ignoring its shape throughout the day. In contrast, our approach focuses on the shape of the consumption pattern revealing substantial differences between articles. Thus, using the total number of pageviews (volume) as a single relevance metric will likely miss nuances about how these articles are useful to the reader. Complementary metrics capturing the usage of articles, such as the shape, describe important properties that might act as confounding factors, and other studies should consider them in stratified analyses. In addition, they could also be directly useful for editors providing additional information about the audience of articles, potentially helping to bridge the gap reported in the misalignment between supply and demand of articles (Warncke-Wang et al. 2015).

Customization. These results should be taken into account when building tools to support users in finding and accessing relevant information (e.g. Wikipedia’s search or recommendations such as RelatedArticles⁵) or to customize their online experience. For example, on average offering information about a movie in the morning has less value than in the evening. Similarly, understanding the content that draws more attention at a different time of the day has implications for designing systems to engage potential editors when they are more inclined to contribute to specific content.

Infrastructure optimization. As observed in previous work (Reinoso et al. 2012), these findings are valuable for optimizing the Wikimedia infrastructure. Given the scale of Wikipedia, optimizing data caching and load-balancing to reflect the actual needs of the consumers across space and time can offer significant benefits for the platform’s performance.

Limitations and future work. Our analysis may have some data limitations, although we argue that it is unlikely they would alter our main findings and conclusions. For example, despite the best efforts of the Wikimedia infrastructure team in developing heuristics to detect bots, the data may still contain some automated traffic. Similarly, country and timezone are identified using the IP address of the request, which may be sensitive to the use of VPNs. Other aspects that may influence our findings are unobserved external factors such as biases introduced by search engines. Since the large majority of Wikipedia traffic originates from search engines (Piccardi et al. 2023), search engine results varying with the time of the day would also impact visits to Wikipedia.

One concrete limitation of our study that we aim to address in future work is focusing only on the English edition of Wikipedia. Readers from non-English-speaking countries consuming content from this edition may be a biased population. Future work should extend our study to multiple languages to compare information behaviors worldwide. Also, future work should study how individual-level (as opposed to population-level) information needs change during the day, possibly combined with demographic data, such as age or education (Kulshrestha et al. 2020). In our case study, this was not possible due to privacy constraints.

Another limitation of this study is the potential impact of external factors. The readers’ behavior could be affected by extended temporal cycles, including annual trends or worldwide occurrences, such as the mobility restrictions imposed during the COVID-19 pandemic (Ribeiro et al. 2021). Consequently, further research should explore these patterns over a period extending beyond the four-week duration of this study.

Moreover, future work should explore the shape and strength of temporal rhythms on other platforms beyond Wikipedia, especially on search engines, which tend to be the first resource we turn to when we need information online. In this sense, studying search engines can offer a complementary view to our findings. Ultimately, we hope that our study can help pave the way to better serve the needs of Wikipedia readers and Web users in general.

⁵<https://www.mediawiki.org/wiki/Extension:RelatedArticles>

Ethical considerations. Server logs may contain sensitive information with implications for the privacy of users. In this work, we pay special care to ensure the researchers access only anonymized records, excluding activities of logged-in users and editors that may be linked through public data. Our findings describe aggregated behaviors that represent a minimal risk for privacy violations. We believe the benefit of presenting them outweighs the potential risks.

Acknowledgements

We would like to thank Leila Zia and Brandon Black for insightful discussions and for reviewing an initial draft of this paper. West’s lab is partly supported by grants from Swiss National Science Foundation (200021.185043, TMSGI2.211379), Swiss Data Science Center (P22_08), H2020 (952215), Microsoft Swiss Joint Research Center, and Google, and by generous gifts from Facebook, Google, and Microsoft

References

- Abdullah, S.; Murnane, E. L.; Matthews, M.; Kay, M.; Kientz, J. A.; Gay, G.; and Choudhury, T. 2016. Cognitive rhythms: Unobtrusive and continuous sensing of alertness using a mobile phone. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, 178–189.
- Aledavood, T.; Kivimäki, I.; Lehmann, S.; and Saramäki, J. 2022. Quantifying daily rhythms with non-negative matrix factorization applied to mobile phone data. *Scientific reports*, 12(1): 1–10.
- Althoff, T.; Horvitz, E.; White, R. W.; and Zeitzer, J. 2017. Harnessing the web for population-scale physiological sensing: A case study of sleep and performance. In *Proceedings of the 26th international conference on World Wide Web*, 113–122.
- Breiman, L. 2001. Random forests. *Machine learning*, 45(1): 5–32.
- Chalmers, D.; Fleming, S.; Wakeman, I.; and Watson, D. 2011. Rhythms in twitter. In *2011 IEEE Third International Conference on Privacy, Security, Risk and Trust and 2011 IEEE Third International Conference on Social Computing*, 1409–1414. IEEE.
- Chi, E. H.; Pirolli, P.; Chen, K.; and Pitkow, J. 2001. Using information scent to model user information needs and actions and the Web. In *Proc. SIGCHI Conference on Human Factors in Computing Systems (CHI)*.
- Dimitrov, D.; Lemmerich, F.; Flöck, F.; and Strohmaier, M. 2019. Different topic, different traffic: How search and navigation interplay on wikipedia. *The Journal of Web Science*.
- Dodds, P. S.; Harris, K. D.; Kloumann, I. M.; Bliss, C. A.; and Danforth, C. M. 2011. Temporal patterns of happiness and information in a global social network: Hedonometrics and Twitter. *PloS one*, 6(12): e26752.
- Doryab, A.; Dey, A. K.; Kao, G.; and Low, C. 2019. Modeling biobehavioral rhythms with passive sensing in the wild: a case study to predict readmission risk after pancreatic surgery. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 3(1): 1–21.
- Dzogang, F.; Lansdall-Welfare, T.; and Cristianini, N. 2016. Seasonal fluctuations in collective mood revealed by Wikipedia searches and Twitter posts. In *2016 IEEE 16th International Conference on Data Mining Workshops (icdmw)*, 931–937. IEEE.
- Fisher, A.; Rudin, C.; and Dominici, F. 2019. All Models are Wrong, but Many are Useful: Learning a Variable’s Importance by Studying an Entire Class of Prediction Models Simultaneously. *J. Mach. Learn. Res.*, 20(177): 1–81.
- Gildersleve, P.; and Yasseri, T. 2018. Inspiration, captivation, and misdirection: Emergent properties in networks of online navigation. In *Complex Networks IX*, 271–282.
- Golder, S. A.; and Macy, M. W. 2011. Diurnal and seasonal mood vary with work, sleep, and daylength across diverse cultures. *Science*, 333(6051): 1878–1881.
- Gullo, K.; Berger, J.; Etkin, J.; and Bollinger, B. 2019. Does time of day affect variety-seeking? *Journal of Consumer research*, 46(1): 20–35.
- Halfaker, A.; and Geiger, R. 2019. ORES: Lowering Barriers with Participatory Machine Learning in Wikipedia. In *Proc. Human-Computer Interaction (HCI)*.
- Heggli, O. A.; Stupacher, J.; and Vuust, P. 2021. Diurnal fluctuations in musical preference. *Royal Society open science*, 8(11): 210885.
- Helic, D. 2012. Analyzing user click paths in a Wikipedia navigation game. In *Proc. International Convention MIPRO*.
- Helic, D.; Strohmaier, M.; Granitzer, M.; and Scherer, R. 2013. Models of human navigation in information networks based on decentralized search. In *Proc. ACM Conference on Hypertext and Social Media (HT)*.
- Koopmann, T.; Dallmann, A.; Hettlinger, L.; Niebler, T.; and Hotho, A. 2019. On the Right Track! Analysing and Predicting Navigation Success in Wikipedia. In *Proc. Conference on Hypertext and Social Media (HT)*.
- Kreitzman, L.; and Foster, R. 2011. *The rhythms of life: The biological clocks that control the daily lives of every living thing*. Profile books.
- Kulshrestha, J.; Oliveira, M.; Karacalik, O.; Bonnay, D.; and Wagner, C. 2020. Web Routineness and Limits of Predictability: Investigating Demographic and Behavioral Differences Using Web Tracking Data. *arXiv preprint arXiv:2012.15112*.
- Kumar Swain, R.; and Kumar Pati, A. 2020. Habitual daily ‘Good Morning’ message senders reveal the status of their own circadian clock. *Biological Rhythm Research*, 51(5): 735–746.
- Lamprecht, D.; Lerman, K.; Helic, D.; and Strohmaier, M. 2017. How the structure of Wikipedia articles influences user navigation. *New Review of Hypermedia and Multimedia*, 23(1): 29–50.
- Lehmann, J.; Müller-Birn, C.; Laniado, D.; Lalmas, M.; and Kaltenbrunner, A. 2014. Reader Preferences and Behavior on Wikipedia. In *Proc. Conference on Hypertext and Social Media (HT)*.

- Lemmerich, F.; Sáez-Trumper, D.; West, R.; and Zia, L. 2019. Why the world reads Wikipedia: Beyond English speakers. In *Proc. International Conference on Web Search and Data Mining (WSDM)*.
- Leypunskiy, E.; Kıcıman, E.; Shah, M.; Walch, O. J.; Rzhetsky, A.; Dinner, A. R.; and Rust, M. J. 2018. Geographically resolved rhythms in twitter use reveal social pressures on daily activity patterns. *Current Biology*, 28(23): 3763–3775.
- Lydon-Staley, D. M.; Zhou, D.; Blevins, A. S.; Zurn, P.; and Bassett, D. S. 2021. Hunters, busybodies and the knowledge network building associated with deprivation curiosity. *Nature Human Behaviour*, 5(3): 327–336.
- Machlup, F. 1983. The study of information: Interdisciplinary messages.
- Maggio, L. A.; Steinberg, R. M.; Piccardi, T.; and Willinsky, J. M. 2020. Meta-Research: Reader engagement with medical content on Wikipedia. *Elife*, 9: e52426.
- Mangel, M.; Satterthwaite, W.; Pirolli, P.; Suh, B.; and Zhang, Y. 2013. Invasion biology and the success of social collaboration networks, with application to Wikipedia. *Israel Journal of Ecology and Evolution*, 59(1): 17–26.
- Mark, G.; Wang, Y. D. o. I.; Niiya, M.; and Reich, S. 2016. Sleep debt in student life: Online attention focus, Facebook, and mood. In *Proceedings of the 2016 CHI conference on human factors in computing systems*, 5517–5528.
- McInnes, L.; Healy, J.; and Melville, J. 2018. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*.
- Murnane, E. L.; Abdullah, S.; Matthews, M.; Choudhury, T.; and Gay, G. 2015. Social (media) jet lag: How usage of social technology can modulate and reflect circadian rhythms. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, 843–854.
- Park, M.; Thom, J.; Mennicken, S.; Cramer, H.; and Macy, M. 2019. Global music streaming data reveal diurnal and seasonal patterns of affective preference. *Nature human behaviour*, 3(3): 230–236.
- Piccardi, T.; Gerlach, M.; Arora, A.; and West, R. 2023. A large-scale characterization of how readers browse Wikipedia. *ACM Transactions on the Web*, 17(2): 1–22.
- Piccardi, T.; Gerlach, M.; and West, R. 2022. Going down the Wikipedia Rabbit Hole: Characterizing the Long Tail of Reading Sessions. In *Proc. International World Wide Web Conference (WWW) - Companion*.
- Piccardi, T.; Redi, M.; Colavizza, G.; and West, R. 2020. Quantifying engagement with citations on Wikipedia. In *Proc. International World Wide Web Conference (WWW)*.
- Piccardi, T.; Redi, M.; Colavizza, G.; and West, R. 2021. On the Value of Wikipedia as a Gateway to the Web. In *Proc. International World Wide Web Conference (WWW)*.
- Pintér, G.; and Felde, I. 2022. Awakening City: Traces of the Circadian Rhythm within the Mobile Phone Network Data. *Information*, 13(3): 114.
- Pirolli, P.; and Card, S. 1999. Information foraging. *Psychological review*, 106(4): 643.
- Rama, D.; Piccardi, T.; Redi, M.; and Schifanella, R. 2021. A Large Scale Study of Reader Interactions with Images on Wikipedia. *EPJ Data Science*.
- Ratkiewicz, J.; Fortunato, S.; Flammini, A.; Menczer, F.; and Vespignani, A. 2010. Characterizing and Modeling the Dynamics of Online Popularity. *Physical Review Letters*, 105(15).
- Reinoso, A. J.; Munoz-Mansilla, R.; Herraiz, I.; and Ortega, F. 2012. Characterization of the Wikipedia traffic. In *ICIW 2012: Seventh International Conference on Internet and Web Applications and Services*, 156–162.
- Ribeiro, M. H.; Gligorić, K.; Peyrard, M.; Lemmerich, F.; Strohmaier, M.; and West, R. 2021. Sudden attention shifts on wikipedia during the COVID-19 crisis. In *Proc. International Conference on Web and Social Media (ICWSM)*.
- Rotman, D.; Vieweg, S.; Yardi, S.; Chi, E.; Preece, J.; Shneiderman, B.; Pirolli, P.; and Glaisyer, T. 2011. From slacktivism to activism: participatory culture in the age of social media. In *CHI'11 Extended Abstracts on Human Factors in Computing Systems*.
- Saha, K.; Liu, Y.; Vincent, N.; Chowdhury, F. A.; Neves, L.; Shah, N.; and Bos, M. W. 2021. Advertiming matters: Examining user ad consumption for effective ad allocations on social media. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–18.
- Savolainen, R. 2006. Time as a context of information seeking. *Library & information science research*, 28(1): 110–127.
- Scaria, A. T.; Philip, R. M.; West, R.; and Leskovec, J. 2014. The Last Click: Why Users Give up Information Network Navigation. In *Proc. International Conference on Web Search and Data Mining (WSDM)*.
- Singer, P.; Helic, D.; Taraghi, B.; and Strohmaier, M. 2014. Detecting memory and structure in human navigation patterns using Markov chain models of varying order. *PLoS ONE*, 9(7): e102070.
- Singer, P.; Lemmerich, F.; West, R.; Zia, L.; Wulczyn, E.; Strohmaier, M.; and Leskovec, J. 2017. Why we read Wikipedia. In *Proc. International World Wide Web Conference*.
- Suh, B.; Hong, L.; Pirolli, P.; and Chi, E. H. 2010. Want to be retweeted? large scale analytics on factors impacting retweet in twitter network. In *2010 IEEE second international conference on social computing*, 177–184. IEEE.
- TeBlunthuis, N.; Bayer, T.; and Vasileva, O. 2019. Dwelling on Wikipedia: Investigating Time Spent by Global Encyclopedia Readers. In *Proc. International Symposium on Open Collaboration (OpenSym)*.
- Trattner, C.; Helic, D.; Singer, P.; and Strohmaier, M. 2012. Exploring the differences and similarities between hierarchical decentralized search and human navigation in information networks. In *Proc. International Conference on Knowledge Management and Knowledge Technologies*.
- Warncke-Wang, M.; Ranjan, V.; Terveen, L.; and Hecht, B. 2015. Misalignment between supply and demand of quality content in peer production communities. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 9, 493–502.

West, R.; and Leskovec, J. 2012a. Automatic Versus Human Navigation in Information Networks. *Proc. International Conference on Web and Social Media (ICWSM)*.

West, R.; and Leskovec, J. 2012b. Human Wayfinding in Information Networks. In *Proc. International World Wide Web Conference (WWW)*.

West, R.; Pineau, J.; and Precup, D. 2009. Wikispeedia: An Online Game for Inferring Semantic Distances between Concepts. In *Proc. International Joint Conference on Artificial Intelligence (IJCAI)*.

Wilson, T. D. 1981. On user studies and information needs. *Journal of documentation*.

Wilson, T. D. 1997. Information behaviour: an interdisciplinary perspective. *Information processing & management*, 33(4): 551–572.

Wilson, T. D. 1999. Models in information behaviour research. *Journal of documentation*.

Wulczyn, E.; West, R.; Zia, L.; and Leskovec, J. 2016. Growing Wikipedia across languages via recommendation. In *Proc. International Conference on World Wide Web*.

Yasseri, T.; Sumi, R.; and Kertész, J. 2012. Circadian patterns of wikipedia editorial activity: A demographic analysis. *PLoS one*, 7(1): e30091.

Yasseri, T.; Sumi, R.; Rung, A.; Kornai, A.; and Kertész, J. 2012. Dynamics of conflicts in Wikipedia. *PLoS ONE*, 7(6): e38869.

Zhou, D.; Patankar, S.; Lydon-Staley, D. M.; Zurn, P.; Gerlach, M.; and Bassett, D. S. 2023a. Architectural styles of curiosity in global Wikipedia mobile app readership. *PsyArXiv preprint*.

Zhou, K.; Constantinides, M.; Quercia, D.; and Šćepanović, S. 2023b. How Circadian Rhythms Extracted From Social Media Relate to Physical Activity and Sleep. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 17, 948–959.

Ethics Checklist

1. For most authors...

- (a) Would answering this research question advance science without violating social contracts, such as violating privacy norms, perpetuating unfair profiling, exacerbating the socio-economic divide, or implying disrespect to societies or cultures? **Yes.**
- (b) Do your main claims in the abstract and introduction accurately reflect the paper’s contributions and scope? **Yes.**
- (c) Do you clarify how the proposed methodological approach is appropriate for the claims made? **Yes. The sections describing the research questions introduce the methods and the relative explanations.**
- (d) Do you clarify what are possible artifacts in the data used, given population-specific distributions? **Yes. See Data.**
- (e) Did you describe the limitations of your work? **Yes. See Discussion.**

(f) Did you discuss any potential negative societal impacts of your work? **Yes. See Ethical considerations. We do not foresee any significant negative impact.**

(g) Did you discuss any potential misuse of your work? **Yes. See Ethical considerations. We do not foresee any significant negative impact.**

(h) Did you describe steps taken to prevent or mitigate potential negative outcomes of the research, such as data and model documentation, data anonymization, responsible release, access control, and the reproducibility of findings? **Yes. See Data.**

(i) Have you read the ethics review guidelines and ensured that your paper conforms to them? **Yes.**

2. Additionally, if your study involves hypotheses testing...

(a) Did you clearly state the assumptions underlying all theoretical results? **N/A**

(b) Have you provided justifications for all theoretical results? **N/A**

(c) Did you discuss competing hypotheses or theories that might challenge or complement your theoretical results? **N/A**

(d) Have you considered alternative mechanisms or explanations that might account for the same outcomes observed in your study? **N/A**

(e) Did you address potential biases or limitations in your theoretical framework? **N/A**

(f) Have you related your theoretical results to the existing literature in social science? **N/A**

(g) Did you discuss the implications of your theoretical results for policy, practice, or further research in the social science domain? **N/A**

3. Additionally, if you are including theoretical proofs...

(a) Did you state the full set of assumptions of all theoretical results? **N/A**

(b) Did you include complete proofs of all theoretical results? **N/A**

4. Additionally, if you ran machine learning experiments...

(a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **N/A**

(b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **N/A**

(c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **N/A**

(d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **N/A**

(e) Do you justify how the proposed evaluation is sufficient and appropriate to the claims made? **N/A**

(f) Do you discuss what is “the cost” of misclassification and fault (in)tolerance? **N/A**

5. Additionally, if you are using existing assets (e.g., code, data, models) or curating/releasing new assets...

- (a) If your work uses existing assets, did you cite the creators? *N/A*
 - (b) Did you mention the license of the assets? *N/A*
 - (c) Did you include any new assets in the supplemental material or as a URL? *N/A*
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? *N/A*.
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? *N/A*.
 - (f) If you are curating or releasing new datasets, did you discuss how you intend to make your datasets FAIR? *N/A*.
 - (g) If you are curating or releasing new datasets, did you create a Datasheet for the Dataset)? *N/A*
6. Additionally, if you used crowdsourcing or conducted research with human subjects...
- (a) Did you include the full text of instructions given to participants and screenshots? *N/A*
 - (b) Did you describe any potential participant risks, with mentions of Institutional Review Board (IRB) approvals? *N/A*
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? *N/A*
 - (d) Did you discuss how data is stored, shared, and de-identified? [Yes. See Data.](#)