

Temporal Rhythms of Information Consumption on Wikipedia

Tiziano Piccardi
EPFL

Martin Gerlach
Wikimedia Foundation

Robert West
EPFL

Abstract

The time of the day regulates our daily rhythm and what information we access. This study uses billion timezone-corrected pageload events to investigate regularities in daily consumption patterns of Wikipedia articles. We found an association between the readers' consumption behavior and the article's topic, country of access, and the device used. These findings emphasize Wikipedia's role as a rich knowledge base that fulfills multiple dynamic information needs, with possible implications for designing flexible systems to address time-based contextual needs.

Keywords: Wikipedia, Logs Analysis, Information Seeking, Readers Behavior, Daily Rhythms

Introduction

Wikipedia is the largest encyclopedia ever created and one of the largest platforms for open knowledge. The English edition alone contains more than 6.5M articles, growing with a rate of 17k new articles each month¹. Given this large pool of articles, Wikipedia covers a wide variety of topics and offers content to fulfill many diverse information needs. These differences in information needs bring readers to visit Wikipedia for reasons that vary from completing a school/work assignment to boredom (Singer et al., 2017). Previous studies (Piccardi et al., 2023) investigated how readers consume the content on Wikipedia. However, they overlooked a crucial dimension of information needs: the temporal daily rhythms. Although *time* is recognized in information science as a crucial contextual factor that drives human information seeking (Savolainen, 2006), it is often poorly investigated. The diurnal cycle is influenced by many aspects, including our circadian rhythm and environmental aspects that affect the information we need at different times. For example, during an average evening, we may be more inclined to look for information about a TV show than a math theorem. Naturally, given the substantial amount of time we spend browsing the Web, this tendency extends

to online attention, including the content we consume on Wikipedia.

Thus, in this work, we characterize how information consumption in Wikipedia depends on the time of the day and how time interacts with other contextual properties, such as article topics and reader country. We complement previous studies on the consumption and popularity of Wikipedia content by –for the first time– analyzing the access log with timestamps converted into local time to remove the effect of different timezones. Our study demonstrates how readers' demand for the type of information varies substantially throughout the day, and it helps to characterize the population of readers and their needs. These findings serve as the first step toward the design of effective systems to better serve these needs.

Specifically, we address the following research questions:

- rq1** What are the typical shapes of Wikipedia consumption rhythms?
- rq2** Which factors influence the rhythms of Wikipedia consumption?

Data

Our study relies on the access logs collected over four weeks –from March 1st to 28th, 2021– on the servers of the English edition of Wikipedia. These logs offer metadata such as the article loaded and the approximate geolocation. To prepare the data for this study, we select the requests relative to article views, which we anonymize by removing sensitive information as described in previous work (Piccardi et al., 2023). Then, we convert all the timestamps into local time using the timezone information, and we retain 3.45B pageloads associated with 6.3M articles. We represent the number of pageloads of article a for each hour h of the day (averaged over 28 days) by $n_{a,h}$, i.e. each article is represented as a time series with $h = 0, \dots, 23$.

Method

We define the normalized and timezone-corrected consumption pattern of all Wikipedia as the *baseline rhythm*, and it describes the expected access volume variation

¹<https://en.wikipedia.org/wiki/Wikipedia:Statistics>




Figure 1: Four principal components of the daily patterns—obtained from the U matrix of the SVD decomposition—describing together 73.6% of the variance.

over time. We aim to obtain a temporal signature of each article by removing the baseline rhythm and focusing on how their consumption pattern diverges from this global average. We calculate the divergence by computing the per-hour ratio between each article’s time series $n_{a,h}$ (additionally normalized, so each time series sums to 1), and the baseline rhythm $N_h = \sum_a n_{a,h}$.

Fig. 2 summarizes these steps with two articles associated with different topics (STEM and MEDIA). Fig. 2a shows the shape of the curves after the normalization, and Fig. 2b shows their divergence from the baseline rhythm for the same articles.

Results

To approach $rQ1$, we extract the principal components of the matrix describing the normalized 24-dimension time series using SVD. Fig. 1 shows the first principal components in the time domain obtained from the matrix V . Note that each component can contribute positively or negatively to the article pattern. We found that the first component models the day vs. night consumption behavior, capturing 39.1% of the variance. To explore if articles are organized into different prototypical consumption groups, we represent the access pattern with the value of their four principal components and use k -means to cluster the patterns. Interestingly, the articles cannot easily be separated into distinct, discrete groups. Manual inspection of the patterns reduced to two dimensions (UMAP) shows that the consumption behavior of the articles is distributed on a continuum (Fig. 3a), but with a tendency to group by topics (Fig. 3b).

To approach $rQ2$, we decompose each article’s time series by country and access method. Then, we normalize each temporal pattern, remove the baseline rhythm, explode the time series into 24 samples with an explicit feature indicating the hour of the day, and complement with the relative binary vector representing the predicted article’s topics obtained from ORES². The obtained model fits the data with an R^2 of 0.181.

Sorting the 24 coefficients representing the interactions allows us to characterize the temporal shape of each

factor. Fig. 4 shows the temporal shapes of the topics organized in the five top-level groups. The plot highlights 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 *Biographies* 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-owl readers. 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 to articles about *Food* registers an increase during common meal times. Next, Fig. 5 shows the interaction coefficients of the countries with time. Some countries share similar daily patterns. Readers from the United States and Germany tend to consume Wikipedia more than average in the early morning, differently from readers from India and Ireland, who during the same hours consume less content than average. Some countries, like Italy and Spain, reveal shared habits, such as reducing information consumption around noon, possibly associated with lunchtime. Finally, we observe (Fig. 5) that access from desktop devices is above the global average in the central hours of the day, between 9 AM and 5 PM.

Conclusion

In this study, we have provided the first overview of how our daily rhythms affect how we consume information in one of the largest open knowledge platforms, Wikipedia. We hope our study provides a step towards operationalizing information consumption to better serve the needs of Wikipedia readers and web users in general.

References

- [Piccardi et al.2023] Tiziano Piccardi, Martin Gerlach, Akhil Arora, and Robert West. 2023. A large-scale characterization of how readers browse wikipedia. *ACM Trans. Web*, jan.
- [Savolainen2006] Reijo Savolainen. 2006. Time as a context of information seeking. *Library & information science research*, 28(1):110–127.
- [Singer et al.2017] Philipp Singer, Florian Lemmerich, Robert West, Leila Zia, Ellery Wulczyn, Markus Strohmaier, and Jure Leskovec. 2017. Why we read Wikipedia. In *Proc. International World Wide Web Conference (WWW)*.

²<https://www.mediawiki.org/wiki/ORES>




Figure 2: Example of daily pattern of two articles about different topics – STEM dashed blue curve and Media dotted red curve. (a) Normalised daily time series ($n_{a,h}$). (b) Divergence from the normalised baseline rhythm ($D_{a,h}$).

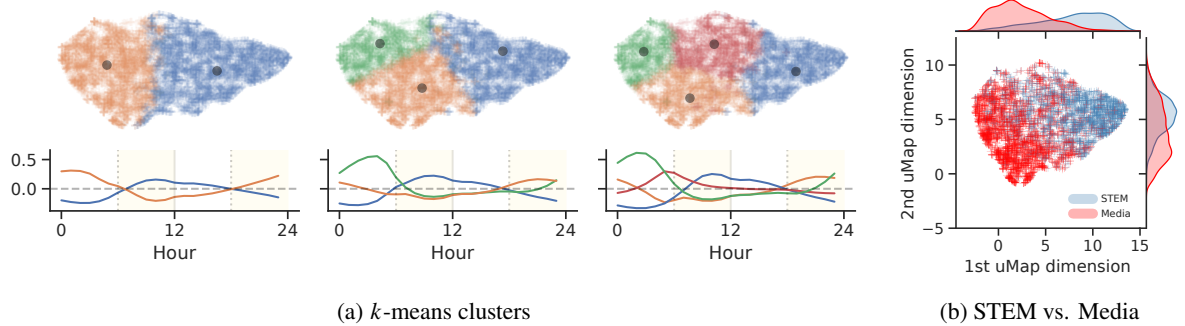


Figure 3: (a) Top: UMAP projection of 10K articles. Clusters obtained with k -means for different values of $k = 2, 3, 4$. Bottom: Shape of the centroids reconstructed from their first 4 principal components. (b) Articles with color assigned based on their topics




Figure 4: Coefficients of the interaction between topics and hours. Background colors represent the high-level classification: STEM (green), Media (yellow), Culture (blue), History & Society (red) and Geographical (grey).




Figure 5: Coefficients of the interaction between country with hours (left, gray background, sorted by the most frequent origin of the requests), and device with hours (right, last plot). Yellow area for the desktop coefficients represents the typical working hour in western countries.