



Between news and history: identifying networked topics of collective attention on Wikipedia

Patrick Gildersleve¹ · Renaud Lambiotte² · Taha Yasseri³

Received: 18 November 2022 / Accepted: 6 June 2023 / Published online: 8 July 2023
© The Author(s) 2023

Abstract

The digital information landscape has introduced a new dimension to understanding how we collectively react to new information and preserve it at the societal level. This, together with the emergence of platforms such as Wikipedia, has challenged traditional views on the relationship between current events and historical accounts of events, with an ever-shrinking divide between “news” and “history”. Wikipedia’s place as the Internet’s primary reference work thus poses the question of how it represents both traditional encyclopaedic knowledge and evolving important news stories. In other words, how is information on and attention towards current events integrated into the existing topical structures of Wikipedia? To address this, we develop a temporal community detection approach towards topic detection that takes into account both short term dynamics of attention as well as long term article network structures. We apply this method to a dataset of one year of current events on Wikipedia to identify clusters of Wikipedia articles related to news events, distinct from those that would be found solely from page view time series correlations or static network structure. We are able to resolve the topics that more strongly reflect unfolding current events vs more established knowledge by the relative importance of collective attention dynamics vs link structures. We also offer important developments by identifying and describing the emergent topics on Wikipedia. This work provides a means of distinguishing how these information and attention clusters are related to Wikipedia’s twin faces of encyclopaedic knowledge and current events—crucial to understanding the production and consumption of knowledge in the digital age.

Keywords Wikipedia · News · Community detection · Collective attention

R. Lambiotte and T. Yasseri have contributed equally to this work.

Extended author information available on the last page of the article

Introduction

Since its founding in 2001, Wikipedia has grown from a simple, often dismissed, ‘Web 2.0’ based dream of shared knowledge to the Web’s primary authoritative information resource for billions of people. Its emergence has overhauled traditional conceptions of the encyclopaedia to become a rapidly updatable real-time record. In addition, it may act as an audience barometer for both historical knowledge and, crucially, current events. Wikipedia is thus an intriguing site of study for news events—reflective of collective audience trends, yet somewhat divorced from the influence of journalistic and content delivery processes.

Wikipedia is a particularly attractive platform to analyse audience reception of current events. Beyond its position as one of the most popular sites on the World Wide Web, the information on Wikipedia is stored in a hyperlink network of articles. Information representation is then a result of both the knowledge structure on Wikipedia and the way its readers navigate that structure. This is unlike a news website, where there may be a single news article, or occasionally a small series of articles towards each event, often disconnected from each other, despite sharing similar subjects. Wikipedia data then facilitates connections and comparisons between different events, as well as to longer term historical knowledge structures. Moreover, using Wikipedia as a data source to study current events fulfils calls for the use of “extra-media data”—information about current events not taken directly from news media [1]. Various works have taken advantage of Wikipedia’s online importance, relevance towards current events, and data availability, using different aspects of the real-time digital encyclopaedia in studying areas such as stock prices [2], film box office performance [3], and disease spread [4].

To holistically study the interaction of current events with Wikipedia, we must first identify what events are represented and how they are manifested. Individuals browse groups of articles relating to events in the news, but what are the various groups that are browsed? Are they consistent between events in forming news topics? And how well do access patterns align with the structure of the information in the encyclopaedia? One must answer these questions to study how different kinds of events are integrated into collective memory, encyclopaedic knowledge, and the historical record.

Previous research has made progress in exploring different modes of data, such as page views, hyperlinks, and edits [5, 6], as well as focusing on specific events or event categories, such as natural disasters or sporting events [7, 8]. However, these efforts have often been limited in scope and/or have not been able to connect directly with news records [9, 10]. Therefore, there is a need for comprehensive techniques that can link specific events, integrate page view attention patterns with hyperlink network structure, and extrapolate individual events to broader topics, all simultaneously.

To address this gap, in this work we present a topic detection model for Wikipedia. Our method leverages both Wikipedia hyperlink network structure and correlations between page view time series, combined with a database of events from Wikipedia’s current events portal, to identify groups of articles that are both well

connected *and* exhibit similar patterns of page views around individual events. These groups of articles are then connected through time to identify the recurrent topics attracting attention on Wikipedia. The resulting ‘Event Reactions’ and ‘Topics of Attention’ encompass news topics as well as attention towards background topics on Wikipedia (and can be resolved as such). These objects are revealing of various interesting features on the nature of knowledge and news recorded on Wikipedia and can act as general purpose event and topic summaries for study in future work. Our work is built around the following research questions.

- **RQ:** How are current events represented in the knowledge structures and access patterns of Wikipedia and its users?
 - RQa: Can we meaningfully and robustly sample the related groups of Wikipedia articles associated with a given news event?
 - RQb: Are these groups of articles from different events related and do they form coherent topics?

We find that our approach yields communities of articles (‘Event Reactions’) that are strongly related according to both hyperlink network structure and correlated page views. When considering Event Reactions’ relationships to each other we find that they form coherent, human-validated, higher-level ‘Topics of Attention’. These topics may be resolved to volatile news topics and stable background topics, representing both facets of Wikipedia as a stable knowledge base and rapidly updating current events record. The topics themselves also qualitatively exhibit a background historical concept space, strong geographical effects, a focus on individuals, and breakout subtopics.

Related work

Historic variation in news and encyclopaedism is based on the discrepancy in information distribution methods. The Internet has narrowed, even brought about a convergence, in the time, format, and audience for these two initially very different knowledge exchange media. Through Wikipedia the encyclopaedia as a format has come to be a highly responsive trove of information on current events, plus readily accessible online news archives (occasionally from citizens’ records) comprehensively chart the events of the 21st century. News and encyclopaedic recording, together with the public’s experience of them, have never been more similar. This convergence has been thoroughly explored [11, 12], and is also reflected in reckonings on its significance to “open-source history” [13], and collective memory [14–21]. Issues of contestation and distortion of this prominent digital record of collective memory are also attracting increasing attention, for example, the cases of widespread administrator-driven far-right historical revisionism on Croatian Wikipedia [22] and distortion of Holocaust history on the English edition [23].

Theoretical work on news media has long had to wrestle with the fact that news outlets themselves, in (justifiably) selectively covering events, are not necessarily representative of current events more widely or the audiences' views on them. Karl Erik Rosengren calls for studies that incorporate “extra-media data”—information about current events not taken directly from news media—to help address this issue [1]. In the Internet age, we may now turn to the large scale tracking of user access patterns and active constructions of repositories of knowledge for a fresh perspective on this extra-media data.

Wikipedia is of course not a news website. It is, however, used as a secondary information resource, which individuals use to further research and contribute towards topics they have encountered through other news media. According to respondents to Wikimedia surveys, 13% of readers visit the site directly because of current events, and a further 30% visit due to wider media coverage¹ [24]. Perceptions around the reliability of Wikipedia have only improved since its early days, teachers' and lecturers' warnings have grown increasingly futile, with various studies confirming its overall reliability on a range of subjects [25–29]. From bar wagers to academic papers [30], its authority as the unofficial arbiter of social facts is undeniable. As such, Wikipedia is an appealing representative site for predicting external collective behaviour, both on and offline, from search trends [31], to stock and cryptocurrency prices [2, 32], film box office performance [3], city tourism numbers [33], disease spread [4], and election results [34, 35].

Moreover, plenty of web companies rely on Wikipedia's content in powering their own services. For example, platforms such as Facebook, Google, YouTube, Twitter, Amazon Alexa, and Apple Siri use Wikipedia in producing their own knowledge graphs, informing automated search results and infoboxes, verifying of notable persons, directing their users to authoritative sources on issues of conspiracies and misinformation, and in building powerful large language models [36–41]. The collective telling of events from Wikipedia, and the aggregate user behaviour in browsing it, is thus emblematic of the kind of audience-centric extra-media data required for studying news media.

Studying the news through Wikipedia article data can be considered a form of content analysis, a common technique in communication and journalism studies. Of particular interest is the task of automatic content analysis, whereby topics, trends, agenda, etc. are analysed across large corpora of news stories where manual coding is not feasible. Three main forms of this are identified in [42] and [43]; rudimentary dictionary based methods (e.g. [44]), supervised machine learning approaches (e.g. [45]), and unsupervised machine learning approaches (e.g. [46]). A novel unsupervised approach in which individual news stories are clustered into ‘news story chains’ according to textual similarity (particularly with novel words), grouping individual articles and their follow-ups into single entities is offered by [47].

A key feature of Wikipedia that separates it from traditional encyclopaedias and news media (as well as in truth much of the modern social web) is the fact that content on Wikipedia exists in a single hyperlinked network of articles. What is

¹ No noted percentage for usage of Wikipedia as a news source.

constructed by different groups of editors, as well as how it is built, better informs us of the information itself. Prior literature has explored how these cultural differences manifest on the site [48], how the network itself drives rich collaborative environments [49], as well as how related knowledge graphs can be used for automated fact checking [50]. Previous literature typically concerns itself with events that have dedicated articles [51, 52]. However, news events can also be documented within one already existing article or across several different articles. For example, stories about one public figure making newsworthy comments about another may be separately recorded on their respective articles.

Work on the concept of spillovers is also relevant here. Even in cases where content in linked articles is not directly relevant to the focus of some exogenous shock, attention in page views and edits may still ‘spillover’ to a selection of neighbours. This has been studied in the context of current events [18, 53], but also the effects of articles being featured on the Wikipedia main page [54], or editing campaigns [55]. Analysing the dynamics of news events in the context of their links to related topics thus clearly necessitates a network based approach.

Several papers cover methods for the detection and summarisation of news events using activity on the platform, frequently relying on bursts in edits [5] or page views [6, 56], or alternatively dedicated databases for particular event categories [8]. Collecting and comparing events from a common category can be informative in answering how event dynamics vary according to category specific parameters. However, further work is needed in comparing across categories. Efforts have also been made to enhance news event detection by combining Wikipedia with other forms of social media acting as a source [51] or filter [57] for events. Bursts of attention towards a topic seem like a natural way of selecting news events but clearly do not cover the full expressible range of event dynamics. Case studies considering edit and page view dynamics [7] or work studying breaking news events more generally [58] look to address how users’ collective attention towards articles situated in a wider topic network evolves, and how it can drive the collaborative editing activities that shape the content and structure of knowledge on the website. News events are rapidly covered on Wikipedia, yet they can have a lasting impact on article content and network structure, incrementally contributing to the collective knowledge base of Wikipedia.

In doing so, individual events shape and integrate into the wider topics represented on Wikipedia. Work that approaches the task of identifying and summarising these topics can use semantic information [59, 60], the article network structure [56, 61], category tags [62, 63], and page view patterns [9, 10, 56]. Several of these approaches are also language agnostic [9, 10, 56, 64], or even multi-lingual [59, 60]. Most notable of these approaches is that of [9, 10, 56] whose language agnostic community detection model incorporates correlated page views together with article network structure. However, in cases where larger datasets are used (such as the topic-level analyses) it is frequently the case that properties such as page views are studied independent of any explanatory description, with any detected interesting features such as peaks later being ascribed meaning by the researcher(s) (likely to some external event) [65]. There is an important distinction between starting from the point of current events and understanding their dynamics, rather than observing

particular dynamics (such as bursts and anomalies) and later attempting to relate them to news events. Firstly, sampling-wise we may only select for particular dynamics when adopting the latter approach (as already touched upon). Secondly, immediately linking to news events assists in later stages in the interpretation of results.

Wikipedia data

The three primary classes of Wikipedia data used are information on events that occur from the Wikipedia Current Events Portal [66], data for the article network of Wikipedia (i.e., the article names and what hyperlinks exist between them), and time series data for the daily page views to each article. Supplementary data on Wikipedia redirects is also used. Further detail on data and how it is obtained may be found in Appendix A. All code and data is available through the WikiNewsTopics GitHub repository.²

Current events portal

The Wikipedia Current Events Portal is a daily archive of events as recorded by Wikipedia editors. Whilst records are in English, coverage is (nominally) of global events of international interest. Events are sorted in 10 categories and are recorded with a summary sentence with links to relevant articles. A partial snapshot is displayed in Fig. 1. We scrape the page to sample a full year of events from 1st December 2017 to 30th November 2018. Initial data gathered for each event includes the date, category, full text description, and the linked Wikipedia articles (henceforth referred to as “core articles”) in each description. For example, in Fig. 1, the final event on 01/04/2017 in the ‘Disasters and Accidents’ category is described as “Authorities cannot contact the *South Korean* cargo freighter *Stellar Daisy*. It is believed that the ship sunk off the coast of *Uruguay*”. We extract the linked pages *South Korea* (displayed text does not have to match article title), *Stellar Daisy*, and *Uruguay* as the core articles for this event. In total, 7919 events are gathered from this year-long period.

Clickstream networks

For the Wikipedia network data, we use dumps from the Wikipedia Clickstream [67]. The Wikipedia Clickstream contains monthly aggregated counts for the number of times links are accessed on Wikipedia, and crucially where from, in (referrer, resource)—equivalently (source, target)—pairs. Only hyperlinks which are clicked more than 10 times in a month are included in the dataset. We only include hyperlinks between Wikipedia articles, excluding links from external

² <https://github.com/pgilders/WikiNewsNetwork-01-WikiNewsTopics>.

Portal:Current events [edit source]

This is an archived version of Wikipedia's Current events Portal from April 2017.

April 1, 2017 (Saturday) edit history watch

Armed conflicts and attacks

- Military intervention against ISIL
 - An Iraqi airstrike near Al Qaim, Anbar province, Iraq, kills Ayad al-Jumaili, believed to be ISIL's second-in-command. *(Al Jazeera)*

Disasters and accidents

- Cyclone Debbie
 - Consequences of heavy rain from Cyclone Debbie kill at least three people and force 20,000 others to leave their homes in New South Wales and Queensland. *(News Limited)*, *(ABC)*
 - Police search for four missing people in southeast Queensland as the Logan River reaches record levels. *(AAP via Yahoo News Australia)*
- A landslide hits the Indonesian island of Java and leaves more than two dozen people missing. *(DPA via News Limited)*
- 2017 Putumayo landslide
 - A landslide in Colombia's southwestern border department of Putumayo sends mud and debris crashing onto houses killing over 250 people and injuring at least 400 others. In addition, 200 people are missing. *(Hindustan Times)*, *(Reuters)*
- An explosion occurs at a carnival in Villepinte, Seine-Saint-Denis, France, injuring at least 18 people. *(The Independent)*, *(BBC)*
- Authorities cannot contact the South Korean cargo freighter *Stellar Daisy*. It is believed that the ship sunk off the coast of Uruguay. *(SBS Australia)*

Politics and elections

- 2017 dissolution of Venezuelan National Assembly

<< **April 2017** >>

S	M	T	W	T	F	S
						1
2	3	4	5	6	7	8
9	10	11	12	13	14	15
16	17	18	19	20	21	22
23	24	25	26	27	28	29
30						

Ongoing events

Disasters [edit source]

- Atlantic hurricane season
- Oklahoma earthquake swarms
- Pacific typhoon season

Law [edit source]

- Philippine Drug War

Political [edit source]

- European migrant crisis (timeline)
- French Guiana unrest
- Romanian protests
- South China Sea disputes
- United States immigration ban
- Venezuelan constitutional crisis
 - Timeline of Venezuelan protests

More details on ongoing conflicts below

edit sidebar

Fig. 1 A snapshot of the Wikipedia current events portal. Live version available at https://en.wikipedia.org/wiki/Portal:Current_events

to Wikipedia and from Wikipedia's Main Page. Essentially, this represents an edgelist that forms a directed, weighted network of monthly navigation between Wikipedia articles. Clickstream data for the English Wikipedia was downloaded for November 2017–December 2018, allowing a one month buffer for studying events at the start and end of the time period of study.

Page view time series

Page view data in hourly granularity for all articles is downloaded from the Wikimedia data dumps [68]. We identify the networks of articles linked to the entries in the Current Events Portal for which page view time series are required (more details in Sect. 4.1), and process the raw compressed time series data to more accessible HDF5 format. This data were downloaded for the period November 2017–December 2018.

Redirects

Wikipedia article redirects are not resolved in all data sources. However, we must account for their important role in Wikipedia's structure and the shaping of traffic [69]. Wikimedia API calls for redirects [70] were used to create a mapping using (1) what 'correct' article title they redirect to (if necessary) (2) all other names that redirect to the article. Page views for individual articles were then calculated by summing those for the groups of their redirects. When mapping the page view data to the articles in the clickstream data this guarantees correct correspondence. The different forms of data are summarised in Table 1.

Methods

Building Event Networks

Here we detail the exact pipeline by which we generate a network of related articles and associated page view time series related to each event (the 'Event Networks'), analyse these for communities of articles representing distinct content and dynamic based 'Event Reactions', and finally cluster these communities (the 'Event Reactions') based on overlapping constituent Wikipedia articles to identify 'Topics of Attention'. These concepts are the key levels of analysis in this work, and are more clearly defined as follows:

- **Event Network:** The hyperlink network of Wikipedia articles and associated page view time series related to a particular news event.
- **Event Reaction:** A community of articles within a single Event Network that are relatively strongly linked and receive correlated patterns of page views.
- **Topic of Attention:** A cluster of Event Reactions from different events, grouped according to common constituent Wikipedia articles.

Firstly, entries from the Current Events Portal must be related to networks of Wikipedia articles and page view data. The process to generate 'Event Networks' runs as follows:

- For each event:
 - Scrape event data from Current Events Portal.
 - Resolve redirects of 'core articles' linked in event description.
 - Use clickstream data to create network of all articles that link to, and are linked to by, as well as all links between these articles over a window of 61 days centred on the recorded event date. Edge weights are a weighted average of the monthly click totals (weights based on fraction of 61 day window in each month).
 - Keep all edges with weight > 100 (i.e. removing edges clicked infrequently), remove any isolates. This is done primarily due to computational speed and

Table 1 A summary of the Wikipedia data used. Further detail available in Appendix A

	Timespan	Resolution	Source	Redirects (raw data)	Notes
Current Events Portal	12/17-11/18	Daily	Scraped	Unresolved	Categorised by news category
Clickstream Networks	11/17-12/18	Monthly	Data Dumps	Resolved	10 click threshold for each link (edge)
Page View Time Series	11/17-12/18	Hourly	Data Dumps	Unresolved	Separate mobile & desktop page views

memory limitations regarding the sparsity of graphs during later community detection stages.

- Collect all article names in the networks, with all redirects, and the period of time they are in the news and require page view data for.
- Process the page view data, keeping data for all required article names, with redirects, over the required time periods.
- Assign 61-day time series (30 days before/after event date) to each event for each article in the respective networks.

The Event Networks encapsulate the network structure and page view dynamics in both anticipation and response to current events, however, they are not a wholly satisfactory description when attempting to generate summary statistics for each event. It is not the case, given the breadth of pages included in each network, that all their articles will exhibit the same signals for page views or edits, many simply being unrelated in the context of the news event. As such, simple averaging techniques for network level features will likely wash out any useful information.

From Event Networks to Event Reactions

One might think that the issues with the Event Networks are a result of the sampling strategy. In abstract terms, there may well be one true signal for each news event, yet it is obfuscated by the noise of less related pages picked up in the network, or concurrent news events involving the same pages, and that the solution is simply some filter or averaging process. We argue that on the contrary it is the very nature of current events that when studying the underlying constituent concepts there may be a variety of responses. These may be due to longer term effects from historical events, structural effects from related information, as well as associations with other current events. This could lead to a variety of different page view patterns. This may seem trivial, but it often does not explicitly emerge in research where the objects of study in focus are specific individual hashtags, news articles, YouTube videos, etc.

The constituent Wikipedia articles relating to individual events exhibit a variety of different dynamics tied to historical, structural, and concurrent news effects. We thus propose a method to separate responses across both content (structure) and attention (dynamics), to identify which groups of articles are both well connected *and* exhibit similar page view time series. Simply taking clusters according to network structure ignores short term associations. On the other hand, simply taking pages with correlated responses ignores context of the related content, and could also introduce spurious associations. This approach takes both factors into account.

The chosen two-stage temporal community detection approach disentangles the different response signals across each Event Network into communities of articles termed ‘Event Reactions’. Each news event is partitioned into a handful of ‘Event Reactions’ across the different subjects represented. These ‘Event Reactions’ are then clustered with those from other news events according to common constituent

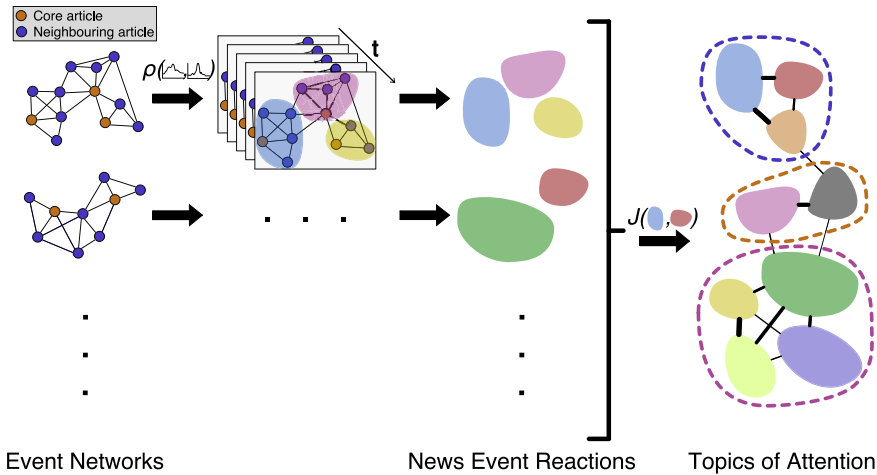


Fig. 2 A schematic of the processing of Event Networks to Event Reactions to Topics of Attention. For each Event Network, edge weights from rolling page view Pearson correlations (ρ) are calculated between article time series. Temporal community detection extracts the Event Reactions which (through Jaccard similarity J) form the higher level-network. We then identify the Topics of Attention with another stage of community detection

articles (via a weighted Jaccard index), to detect broader topics, termed ‘Topics of Attention’. A schematic of the full process is shown in Fig. 2.³

Temporal community detection of signals on knowledge networks

Together with Sect. 4.1, here we tackle **RQa**: *Can we meaningfully and robustly sample the related groups of Wikipedia articles associated with a given news event?* For a given news event i , with associated network of articles $G_i(V, E)$ (nodes representing articles, edges representing hyperlinks between them), there is an associated set of time series for page views towards articles $p_n(t) \forall n \in V$ with T timesteps. To measure similarity in patterns of collective attention towards articles, we calculate Pearson correlations of these time series for all linked nodes over a rolling 7 day window, yielding W_{edges} , which is a $\|V\| \times \|V\| \times (T - 6)$ dimensional array. The full correlation matrix would be represented by W , so that $W_{\text{edges}} = W \circ (A + A^T) / 2$, where A is the (unweighted) adjacency matrix of G_i and \circ denotes the Hadamard product between the matrices. We exploit the information in the sparsity of the hyperlink network to only calculate correlations for a tiny fraction of node combinations, allowing the algorithm to scale well to large networks. The operation also enriches the static network structure with temporal information. Related approaches

³ The reader familiar with topic modelling may find the following analogy useful. Individual Event Networks represent ‘documents’, that are made up of Wikipedia articles, akin to ‘words’, and each news event represents a sample of wider news ‘topics’, that we term ‘Topics of Attention’.

could find use in other domains, e.g., in neuroscience to combine structural network information (DTI) with time series (fMRI) [71, 72].

W_{edges} represents an undirected (note that Pearson correlation is a symmetric measure, which motivates the removal of directedness), weighted temporal network $G'_i(t)(V, E(t))$, on which we perform community detection to identify groups of articles that are both well connected by hyperlinks, and exhibit correlated patterns of page views.

The Leiden algorithm [73] is selected for temporal community detection. This is an extension of the popular Louvain algorithm [74], but addresses an issue whereby communities may be arbitrarily badly connected, it also runs faster than the Louvain algorithm. Rather than standard modularity, the Constant Potts Model [75] is used for the quality function, since it can handle both positive and negative edge weights (which can in principle be observed), the readily interpretable resolution parameter, and the independence of communities from the observed graph/subgraph (particularly important given the articles present are a sample of the much larger Wikipedia article network). To extract information from the temporal network, we further adapt the method proposed in [76], by considering the $T - 6$ layers of the temporal graph, and connecting the same node in successive layers by an interlayer edge with weight $\tau = 1$. By doing so, the temporal network is represented as a static, weighted network where each node appears $T - 6$ times. We optimise its Constant Potts Model with the Leiden algorithm, thereby uncovering communities that are made of nodes at multiple times. Note that this operation requires the determination of the resolution parameter. A search for this parameter with a robustness test on a 100 event sample is carried out in Appendix B, with the resolution being set to $r = 0.25$.

For each Event Network, the obtained partition \mathbf{P}_i is comprised of a handful of communities C_{ij} . Any detected communities which contain at least one of the ‘core articles’ from the descriptive text of the event, and that overlap in time with the day of the event are kept as Event Reactions— R_{ij} . Each of these elements is in effect a building block of wider Topics of Attention. The discrepancy in timescales between fast-paced attention towards news events and the more slowly evolving structure of the Wikipedia article network means these topics are not necessarily reflected in solely the hyperlink structure, or solely through correlated short term page views. In addition, satisfactory temporal community detection on one network for one year over the ≈ 6 million English Wikipedia articles is not computationally feasible. From the 7,919 events, we obtain 7,823 Event Networks with more than 1 node and edge (since a small number of event records do not have a popular associated article), and generate 26,579 Event Reactions.

Community detection comparison

Capturing excess page views

Our objective is to collect as good a selection as possible of Wikipedia articles, representative of a particular event. We expect articles related to some current event will exhibit a heightened level of page views around the time of the event. To describe

the dynamics of attention around the event, we should select the communities containing core articles that on average exhibit excess page views around the time of the event. Communities that do not contain core articles are deemed not relevant to the event—the constituent articles are not structurally connected well enough to the core articles and/or page views do not follow a similar enough pattern. In cases where a community contains a core article but does not exhibit an increase in page views, we can conclude the core article, and the rest of the community, are background articles not directly relevant to the event, and can put them aside when focusing on event page view dynamics. We can then measure how well we have captured the event with the total excess page views towards articles in the identified relevant communities. We can compare how different community detection approaches perform on this measure on each event i , defined as

$$\text{Excess}_i = \sum_{\substack{R_{ij} \in \mathbf{P}_i, \\ \max_{-1 \leq t \leq 1} (\tilde{q}_{ij}(t)) > 3}} \sum_{k \in R_{ij}} \sum_{t'=0}^{t'=6} p_k(t') - \text{median}(p_k(t)). \quad (1)$$

Here R_{ij} refers to an Event Reaction (community containing a core article) in partition \mathbf{P}_i , $\tilde{q}_{ij}(t)$ is the total page views towards articles in Event Reaction j , centred on the median value and scaled to the interquartile range, and $p_k(t)$ is the page view time series of article k in Event Reaction j . Effectively, by the condition on the first summation, we only consider communities that contain a core article, and that have increased overall views around the time of the event. For each of these selected communities, we sum the page views above the median values over the first 7 days for each of the constituent articles.

To make an instructive comparison, we can consider how community detection approaches with just structural hyperlink information, aggregate navigational information, and our method perform on the excess views measure. For each event, we calculate the captured excess views in communities from our temporal approach against that from a simple implementation of the Leiden algorithm on both the static, unweighted, “structural” hyperlink network, as well as against a static, weighted, “navigational” network, where edge weights are set to the weighted average link clicks from the clickstream data (as used in Sect. 4.1). In both cases, resolutions were selected in a similar fashion to as in the temporal approach, detailed in Appendix B. Comparing the results on the captured excess page views, the temporal approach captures at least as many excess page views as the structural approach in 72.4% of events and at least as many excess page views as the navigational approach in 72.8% of events. Taking the ratios of captured excess views ($\text{Excess}_i^{\text{Temp}}/\text{Excess}_i^{\text{Struc}}$ and $\text{Excess}_i^{\text{Temp}}/\text{Excess}_i^{\text{Nav}}$) and considering the geometric mean across all events, the temporal approach captures 1.13 times the excess views of the structural approach and 1.23 times that of the navigational approach. Taking the median, the temporal approach captures 1.04 times the excess views of both the structural and navigational approaches. Our method better captures articles in communities relevant to current events. An example comparing the three approaches on a single event is provided in Fig. 3.

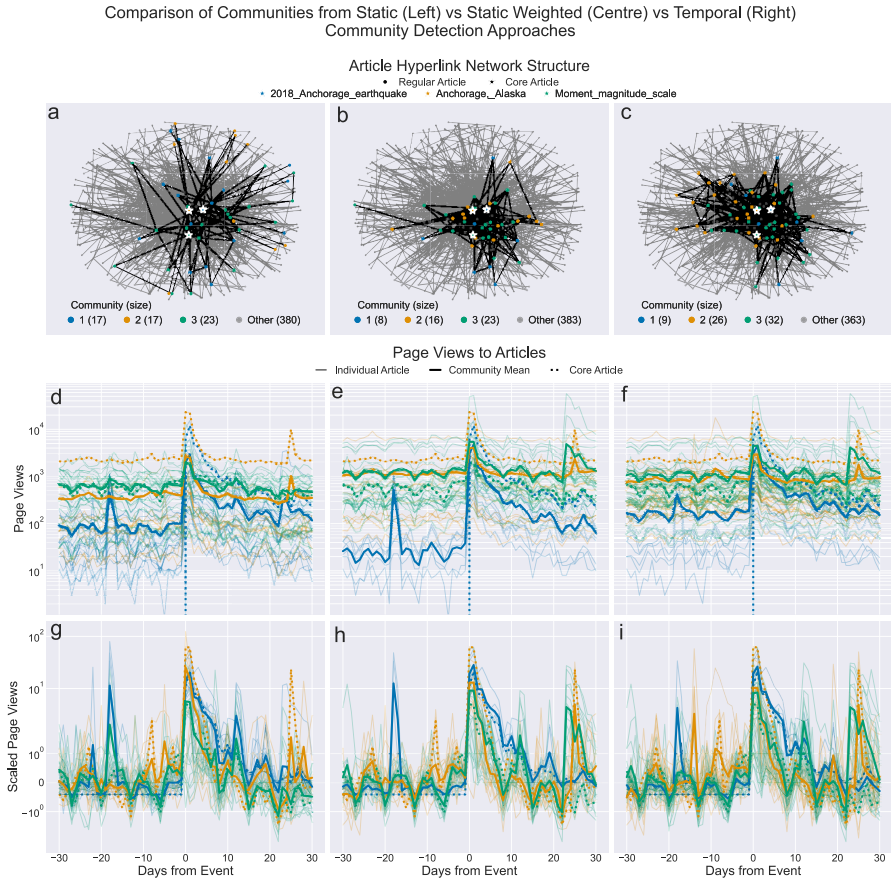


Fig. 3 A comparison between the communities obtained through different methods. The event record in question is “2018/11/30. [[2018 Anchorage earthquake]]: A [[magnitude]] 7.0 earthquake hits Alaska, with the epicenter in [[Anchorage]]. Severe damage is reported.” (core articles indicated by square brackets). In all three approaches there is a community centred around each of the core articles ‘2018 Anchorage earthquake’ (a new article dedicated to the event), ‘Anchorage, Alaska’, and ‘Moment magnitude scale’. The structural, navigational, and temporal approaches capture 263,585, 489,538, and 546,978 excess views, respectively. The absolute page views (in terms of total and mean) increase with the temporal approach (**f** vs **d** & **e**), yet the scaled page view patterns remain similar (**i** vs **g** & **h**). This indicates that a number of additional articles with similar spikes in attention relating to the earthquake have been identified. These articles were not captured in any of the previous static communities. With a static approach, in only taking the articles that are structurally/navigationally close to the “core” articles, we may both miss where attention is being directed by this event and markedly underestimate the amount of attention towards it

Structural similarity

A follow-up task is to compare the makeup of the obtained communities from our combined network structure and page view approach correlations against those generated from a solely network structure approach. If the communities from the

combined approach are no different from those for the network structure baseline, then this indicates attention dynamics in response to any news event have very little effect, and the community is well represented solely by the network. Any ‘disturbance’ by the news event is either minimal, or closely aligns with how information is already represented on Wikipedia. On the other hand, if the communities obtained from each approach are quite different, then the variation in page view dynamics among the articles in the network is important in producing the Event Reactions. Any ‘disturbance’ by the news event is of sufficient magnitude that the association between concepts related to the event is then not well represented by the relatively static network structure on Wikipedia.

The approach for each event is based on comparing the communities already obtained from the temporal network of correlated time series to those we obtain from community detection on a single-layer, unweighted graph representing solely the structure of the article network, without user attention and navigation patterns, (i.e. G_i from Sect. 4.3). For each event i we run community detection on the graph G_i over the same logarithmic range of resolutions $r \in [1.23 \times 10^{-4}, 1]$ from the robustness tests (Appendix B), yielding the partitions \mathbf{P}'_{ir} . Each Event Reaction from the temporal approach (R_{ij}) receives a ‘Structural Similarity’ score s_{ij} . This score is defined as the maximum of the similarities between R_{ij} and each community obtained from the non-temporal approach across all resolutions $C'_{ikr} \in \mathbf{P}'_{ir}$ $1.23 \times 10^{-4} \leq r \leq 1$. Thus,

$$s_{ij} = \max(J_w(R_{ij}, C'_{ikr}) \forall C'_{ikr} \in \mathbf{P}'_{ir} \ 1.23 \times 10^{-4} \leq r \leq 1), \quad (2)$$

where $J_w(x, y)$ is the Jaccard similarity between communities x and y , weighted by the PageRank scores [77] of nodes in the subgraphs x and y [78] (weighting more important articles to the community more highly). This takes into account both the content of each community in terms of Wikipedia articles, and the relative importance of said articles within their respective communities.

The Structural Similarity score describes how dependent the observed community R_{ij} is on variation in short term correlated attention dynamics in an Event Network, compared to the longer term network structure. If all page view time series were uniformly correlated, we would expect $s_{ij} \approx 1$, i.e., all edge weights would be approximately equal, and the community detection is more reliant on the presence/absence of edges. If on the other hand a subset of articles receive strongly correlated page views, uncorrelated with the page views to other articles, we would expect $s_{ij} \approx 0$, i.e., community detection is more dependent on edge weight than simply the presence/absence of an edge. The distribution of s across all Event Reactions is shown in Fig. 4. We observe a range of behaviours; a prominent mode with relatively low structural similarity (i.e. page views are important), a broader intermediate mode (page views have some effect), and finally the sharp mode around $s = 1$ (page views have little to no effect).

Higher-level Topics of Attention

Over all events, we now have a collection of Event Reactions. Many of these will be related through covering different stages of the same continuous event (e.g. different




Fig. 4 Distribution for the structural similarity scores of all Event Reactions

rounds of the FIFA World Cup), or through the re-emergence of events and news topics in time (e.g. updates about the Mueller investigation, or new natural disasters). We now turn to **RQb**: *Are these groups of articles from different events related and do they form coherent topics?* We seek to identify the recurring groups of Wikipedia articles associated with news events—the *topics* that are represented. Event Reactions from different events that are made up of broadly the same collection of Wikipedia articles are representative of a wider concept receiving repeated news exposure. We look to quantify the similarity between Event Reactions and use this to find the more closely related groups that represent Topics of Attention. We construct a higher-level network $H(V, E)$ of all Event Reactions (nodes). Edge weights are set as the weighted Jaccard similarity [78] between the sets of articles of each Event Reaction, weighted by their PageRank centrality in their respective networks [77], indicating similarity in content (and weighting more important articles to the concept more highly). This network contains all recorded instances of Event Reactions in the sample, representing their relation to one another over the course of one year. To identify the Topics of Attention (groups of related Event Reactions) we run a further stage of community detection over this network $H(V, E)$, using the Leiden algorithm with the Constant Potts Model as before. Whilst the nodes in the network represent snapshots of events centred on different points in time, $H(V, E)$ is not a temporal network. The resolution parameter is set at $r = 0.067$, according to the robustness test set out in Appendix C. This process yields a partition of communities that are the Topics of Attention which we go on to label, validate, and explore.

Topic labelling and validation

In line with literature on news values and newsworthiness [79–82], the Topics of Attention were sorted by several features detailed in Table 2, with the top topics across each feature manually labelled. Several of these are based on the constructed

Table 2 A summary of the features with which we sort and examine the Topics of Attention

Feature	Description
Number of associated events	Topics most frequently featured in the news (often background topics)
Prominence	Topics which on average have the largest level of pre-existing attention, median page views (often background topics)
Magnitude	Topics which on average receive the largest increase in page views when in the news
Deviance	Topics which on average receive the largest increase in page views, relative to their prominence, when in the news

time series, W_{ij} for each Event Reaction (R_{ij}). This is a sum of the daily page views to each article k in an Event Reaction ($p_k(t)$), weighted by their PageRank centrality (w_k) to the network they form;

$$W_{ij}(t) = \sum_{k \in R_{ij}} w_k p_k(t). \quad (3)$$

The time series is then centred in time to the max value occurring ± 1 day from the recorded date of the event. For a Topic of Attention A_α with constituent Event Reactions R_i (i acting as an index for the Event Reactions in A_α and no longer referring to a specific event) and their associated time series W_i , the average prominence, magnitude, and deviance of a topic are then accordingly

$$\text{Prominence}_\alpha = \frac{\sum_{W_i \in A_\alpha} \text{median}(W_i(-30, -29, \dots, 0))}{\|A_\alpha\|}, \quad (4)$$

$$\text{Magnitude}_\alpha = \frac{\sum_{W_i \in A_\alpha} W_i(0) - \text{median}(W_i(-30, -29, \dots, 0))}{\|A_\alpha\|}, \quad (5)$$

$$\text{Deviance}_\alpha = \frac{\sum_{W_i \in A_\alpha} \frac{W_i(0) - \text{median}(W_i(-30, -29, \dots, 0))}{\text{median}(W_i(-30, -29, \dots, 0))}}{\|A_\alpha\|}. \quad (6)$$

Intuitively, prominence corresponds to how popular the subject(s) of an event or topic are before a particular event takes place. Magnitude corresponds to the absolute attention towards a given event when it does occur. Finally, deviance corresponds to unexpected event popularity, or how shocking a given event is, relative to its typically fairly unpopular subject matter.

Two coders independently manually labelled a subset of 65 of the Topics of Attention by examining the constituent Wikipedia articles for each topic and the news events most associated with them (five events initially, with the option to see more). The labelling interface is shown in Appendix D. This set of topics were selected by

Table 3 Top topics by certain measures (min 10 events). Colour indicates quartile of structural similarity score, from red=bottom quartile to green=top quartile. Symbol indicates labelling agreement. No symbol: Unanimous Strong, *: Strong-Partial, **: Unanimous Partial, †: No agreement

#	Event Reactions	Prominence	Magnitude	Deviance
0	Countries	FIFA World Cup	British Royal Family	US administration
1	US Politics**	Unknown	FIFA World Cup	California Wildfires 2
2	US Domestic**	British Royal Family	American Football	Russia-Ukraine relations**
3	Israel Palestine conflict	Countries	US administration	Australian Prime Ministers**
4	Syrian Civil War	East Asia**	Tennis	Wildfires
5	Global Cities	US Finance	Indian Politics	Turkey-United States relations**
6	Middle East	Unknown	Brett Kavanaugh Supreme Court nomination**	California Wildfires 1
7	North Korea - South Korea relations*	Trump Family	Malaysian politics	Australian Politics
8	Terrorism in Afghanistan	US Politics**	Thai cave rescue	White Helmets (Syrian civil war)
9	Trump Family	Cryptocurrency	California Wildfires 2	Florida
10	Mueller Investigation	Winter Olympics	Gun violence in the United States	Peruvian Politics
11	Putin & Russian Politics	Big Tech	Malaysian Politics	Israel in Syrian civil war*
12	US political houses	(South)eastern Europe**	Musk & Tesla	Puna Volcano
13	Korean conflict*	US Presidency*	US political houses	Malaysian politics
14	Global Currents	Social Media	Foundation of USA**	Costa Rican politics
15	Latin America	Foundation of USA**	Tropical Storms	Thai cave rescue
16	Yemeni Civil War	Colombian drug trade*	Toy retailers**	Toy retailers**
17	Catalan Independence Movement	US Domestic**	Australian Politics	Yemeni politics
18	Saudi Royal Family	Global Cities	US Gun violence	Iranian Politics
19	UK Politics	US military	Olympics	Georgian politics
20	Tropical Storms	Middle East	NHL	Johnson Family

taking the top 20 topics across each feature in Table 2 (some topics appear more than once across the four features, hence the total less than 80). Each coder was then presented with the combined list of labels and independently tasked with identifying where there was ‘strong agreement’, ‘partial agreement’, or ‘weak/no agreement’ between labels. For the topics, 72.3% of labels were in unanimous strong agreement, 7.7% in strong-partial agreement (i.e., one coder ranked as strong agreement, and one as partial agreement), 15.4% in unanimous partial agreement, and 4.6% in weak/no agreement. The procedure demonstrates validity of the interpretable topics. For the purpose of display in figures and tables, in cases where there was not unanimous strong agreement between coders the first coder’s labels are used.

Results and discussion

Studying the contents of the emergent Topics of Attention in Table 3 reveals various interesting details on current events as recorded on Wikipedia. Identified features include; a background concept space, strong geographical effects (including a heavy Anglosphere/US focussed bias), a focus on individuals, and breakout subtopics.

Background concept space

Several of the top Topics of Attention by number of associated events (Countries, Global Cities, Tropical Storms, etc.) are those of lasting historical context. These topics also typically have high structural similarity—attention towards the topic is correlated with its structural composition on Wikipedia. Whilst the Event Networks are sampled from a current events records, much of the related content is built on and widely considered as part of long established knowledge. This supports the case of news events contributing to longer term narratives.

Strong geographical effects

The Topics of Attention are strongly characterised by geography. Many of the labelled topics are specified by the region they are relevant to. It is also clear that when incorporating the structure of the knowledge graph and attention that many of the most prominent topics are Anglosphere focussed. This is of course partly a consequence of studying the English Wikipedia. However, the current events portal's nominal aim, and that of English Wikipedia as a whole, is to objectively cover global events and knowledge—something it still falls short on. Topics relating to the US and UK are covered with far higher granularity than those relating to other countries. That is, there are several top topics related to the intricacies of US politics, yet other countries typically have all related news summarised within a single topic. Considering the topic labels with some geographic link, 46% are Anglosphere based (primarily US, with a handful UK and Australia based). This is not entirely unsurprising, given prior work on Wikipedia biases [83–86] as well as this work's focus on the English language Wikipedia. Nevertheless, this further validates assertions that rather than being the “sum of all human knowledge” [87], Wikipedia (in its various languages), through its content, structure, and access patterns is highly sensitive to its cultural setting.

Focus on individuals

Several of the labelled topics are focussed on, or strongly feature a powerful individual (e.g. Trump family, Putin & Russian Politics, Musk and Tesla). This points to an audience sensitivity towards people that can be related to or reviled and is reflective of findings on the news values of celebrity/power elite [82].

Breakout subtopics

There are several cases where topics may be strongly related, yet one cluster achieves breakout popularity enough to distinguish itself from the original topic. These could correspond to the well studied phenomena of “media storms” [88], whereby there is intense media focus on a single issue. An example of this is the topic for North Korea–South Korea relations—representing an overview of related articles—and the Korean Conflict—which is the subject of more intense focus around events by the audience as indicated by the differing structural similarity scores. On top of the differing structural similarity scores, the Korean Conflict topic has higher prominence, magnitude, and deviance than the North Korea - South Korea relations topic. Another example is the broader US Politics topic compared to the US political houses or current US administration topic. The former represents stable knowledge attracting attention around the topic and the latter represent new, more unusual combinations of articles more closely associated with current events. This may be

a consequence of the choice of Jaccard similarity for the higher-level graph edge weights, where news events create strong, synchronous deviations from typical page view behaviour across a very small group of articles that are still related to a wider group. Since the number of deviating articles for the individual event is small, the edge weight through Jaccard similarity to other related events with a larger set of articles is also small, leading to it not being included in the Topic of Attention. An alternative similarity metric such as the overlap (Szymkiewicz-Simpson) coefficient could account for this effect, though using this would likely smooth over any break-out clusters, interesting features unto themselves.

Further remarks

The qualitative discussion of results and exploration of content is important in contextualising findings in further works on Wikipedia's coverage of current events, as well as Wikipedia status as an 'independent' data source for news media. Beyond simply being indicators for the notable issues in the news over a year, the detected Topics of Attention and their properties are demonstrative of the assertion that there is a disconnect between the 'editor's' Wikipedia and the 'page viewer's' Wikipedia. The central tension of Wikipedia as both a slow moving encyclopaedic knowledge base and fast moving current events record is displayed in the ways the topics are constructed. In the first mode, the audience's access patterns align with the established structure of knowledge on Wikipedia. The truly interesting mode occurs when the audience's attention does not align with the article network and in effect establishes its own communities of related articles. This collective behaviour is what stretches Wikipedia both towards updating its content and remaining a popular information resource, and away from its traditional encyclopaedic grounding.

There are several limitations to the methods proposed here. Firstly is the issue of this being a single language study. There have been a number of articles on the varying content, coverage, and use of different language Wikipedias based on linguistic, cultural, and national focusses [48, 89–91]. One could contend that this means a single story from one community of people editing and viewing Wikipedia. A strong Anglosphere bias is indeed observed but we see it as the case the English Wikipedia is not the product of, nor the information tool, for a single, large, homogeneous community. [92–94] all observe that certain editors occupy particular roles in lending substantive expertise towards particular categories, whether that be due to identity, education, or other personal interest, and the same would be expected of regular users (to some extent also supported by [24]). In addition, whilst solely applied to English Wikipedia here, the majority of the methods used are language agnostic, and may be swiftly applied to other language Wikipedias, which may be a fruitful avenue to pursue.

The Current Events Portal is clearly not an exhaustive source of news stories, many of which would have no discernible effect on Wikipedia. Explicit editing guidelines state that "Stories added to the main portal page should be of international

interest” [95]. Beyond this restriction, there are a very large number of people who regularly access information related to sports, entertainment, and popular culture, whose news stories are rarely featured on the current events portal. Celebrity deaths, for instance, have their own summary article, rather than residing on the current events portal. The topic map thus does not cover the universe of what might be considered news, and is sensitive to the contents of the news story source, the Wikipedia Current Events Portal, raising the issue of endogeneity. Sampling news events by ‘Wikifying’ [96] alternative sources such as news website RSS feeds would indeed yield a different set of events, though unfortunately due to editorial decisions, we of course arrive at a similar obstacle where there is no objective set of events.

A more thorough comparison between the topic landscape of several different news outlets would be of interest and an immediate application of the developed methods towards agenda-setting research, yet is outside the scope of this work. An argument in favour of choosing the Wikipedia Current Events Portal is that this collaborative recording of news is representative of the collective received importance of events, incorporating what the news recording and accessing communities consider relevant. This, together with time constraints and the simplicity of selecting descriptions already formatted with Wikipedia links, resulted in the decision being made to concentrate on the Current Events Portal.

Conclusion

The encyclopaedic origins of Wikipedia mean it is not set up as a ready-made data source for the study of news events. Typically, long term information for each subject, involving many events, is aggregated in each Wikipedia article, as opposed to aggregating information about multiple subjects at the event level for each event (i.e. like a news website). Events, aside from in rare cases where a particularly notable event justifies its own article, and news topics do not have an established natural representation on Wikipedia. Equally, the broadly consistent, common structured information available to its huge audience, as well as how this audience accesses content on current events, is too appealing to ignore. To take advantage of this, one must establish a framework for event and topic level study. To this end, we have developed an approach for event sampling and topic detection on Wikipedia, with a focus on news topics, that takes into account article network structure, dynamics, and content.

The graph supported correlation network approach towards temporal community detection successfully detects stable Event Reactions, relating both short-term dynamics of attention through page views as well as long-term knowledge structures, thus addressing **RQa**. We have demonstrated its utility in identifying and exploring different Event Reactions, and in their aggregate how they represent Topics of Attention, the objective of **RQb**. These objects of study improve upon those used in prior work for their generality across topics, usage of the knowledge network rather than focus on individual articles, explicit relation to news events, incorporation of

short and long term effects, and lack of reliance on detection through particular attention dynamics. The Topics of Attention on Wikipedia exhibit a background historical concept space, strong geographical effects, a focus on individuals, and breakout subtopics. The Topics of Attention (through their constituent Event Reactions) may also be resolved to volatile news topics and stable background topics. More importantly, they represent both facets of Wikipedia as a stable knowledge base and rapidly updating current events record.

Detecting Topics of Attention using Wikipedia has proven to be a non-trivial task. It is important to encapsulate the contrasting timescales of news and existing knowledge, the many to many relationship between events and topics, together with the corresponding dynamics of attention and memory building. Through this process we gain insight how topics are represented and accessed on Wikipedia and on which events are considered important enough to make it into the encyclopaedic record. Finally, generating Event Reactions, and wider Topics of Attention can enable more detailed event and topic level study in further work. Establishing representations of events and topics, beyond individual articles, on Wikipedia allows us to quantitatively address questions on theories of news media, collective memory, and historical recording in ways not previously possible without this kind of massive audience level data.

Appendix A Wikipedia Data

Here we provide some additional details on the data as described in Sect. 3.

Current events portal

Entries in the Current Events Portal are sorted in 10 categories (Armed conflicts and attacks, Law and crime, Arts and culture, Politics and elections, Business and economy, Science and technology, International relations, Sports, Health and medicine, Disasters and accidents). The criteria for an event's appearance on the page is simply that the community believe it is a significant story, subject to guiding principles such as the length and depth of existing news coverage, though ultimately, decisions are based on consensus of editors on the individual merits of events [95, 97].

Clickstream networks

Visits from popular sources outside Wikipedia are also recorded in the raw clickstream data, although we only include hyperlinks between Wikipedia articles, excluding links from external to Wikipedia and from Wikipedia's Main Page. In addition, in the raw data only links with > 10 clicks over the course of each month are supplied. A comprehensive analysis of all of Wikipedia's historical links would require examining full content dumps or revision histories of every article. The




Fig. 5 Partition stability. Partitions at smaller resolutions tend towards a single community and partitions at larger resolutions tend towards a unique community for each node. There are maxima for each similarity measure around $r = 0.25$

authors of this study believe that clickstream data provides a more efficient means of identifying the most relevant links and outweighs any tradeoff in completeness. The clickstream data is used since firstly it offers a fast, reliable snapshot of the past article network structure of Wikipedia (compared to the time taken wrangling HTML from the complete Wikipedia dumps, or from individual article revisions through API calls), and secondly the weighting allows for a cutoff for spurious links.

Page view time series

The page view data is grouped into monthly datasets with hourly granularity for each Wikimedia project. We focussed on the page views towards articles in the English Wikipedia, which are reported under the prefixes for desktop Wikipedia (en.z), the mobile Wikipedia (en.m), and Wikipedia Zero (en.zero) (a discontinued project where Wikipedia access was available for free in developing countries which in practice accounts for relatively few views).

Redirects

A somewhat minor, yet often overlooked, facet of research on Wikipedia is the issue of page redirects [98]. A redirect is a page which automatically sends visitors to another page with the ‘correct’ title, for example searching for ‘USA’ or clicking a wikilink titled ‘USA’ automatically redirects the user to the page ‘United States’. Redirects are very important since users and editors constantly search for articles, click on links, and edit links with alternative/abbreviated/previous names. In fact, it has been estimated up to 55% of the articles in the main namespace of Wikipedia are ‘redirects’ [69]. To ensure information for identical articles in the data is not being




Fig. 6 Partition stability for static, unweighted community detection. Partitions at smaller resolutions tend towards a single community and partitions at larger resolutions tend towards a unique community for each node. Maximal similarity by AMI motivates selection of $r = 0.030$ going forward




Fig. 7 Partition stability for static, weighted community detection. Partitions at smaller resolutions tend towards a single community and partitions at larger resolutions tend towards a unique community for each node. Maximal similarity by AMI motivates selection of $r = 54.6$ going forward

duplicated, redirects need to be resolved. Redirects are already resolved in the click-stream dataset (though pages may change in name over time e.g. ‘Meghan Markle’ to ‘Meghan, Duchess of Sussex’) but they are unresolved in the page view dataset, i.e. hits on redirect pages are recorded separately from hits for the true page name [99].

Appendix B: Community detection resolution selection

We perform a robustness test during the initial stage of temporal community detection on the Event Networks (G_t) to identify the resolution parameter value which gives stable, meaningful partitions. We test a random sample of 100 events,

repeating the community detection process over a (logarithmic) range of resolution parameters, and compare the similarity in the obtained partitions. Similarity between partition n and $n - 1$ is calculated according to adjusted mutual information and CluSim element-centric similarity [100] and shown in Fig. 5. Based on this test, the resolution parameter for further analysis is set according to the peak similarity around $r = 0.25$.

Similarly, in comparing to methods that only use structural or navigational info (Sect. 4.4), we perform equivalent resolution tests for community detection approaches that use static networks based on the article hyperlink network and the aggregate clickstream navigation counts. Partition similarities for these approaches over a range of resolutions are shown in Figs. 6 and 7. For further analysis with these approaches we then choose to use resolutions of 0.030 and 54.6 respectively.

Appendix C: Higher-level network community detection

The approach of comparing partitions across a range of resolutions in B is repeated for community detection on the full higher-level network H and shown in Fig. 8. In this case, the partition from the maxima around $r = 0.067$ is chosen for further analysis.

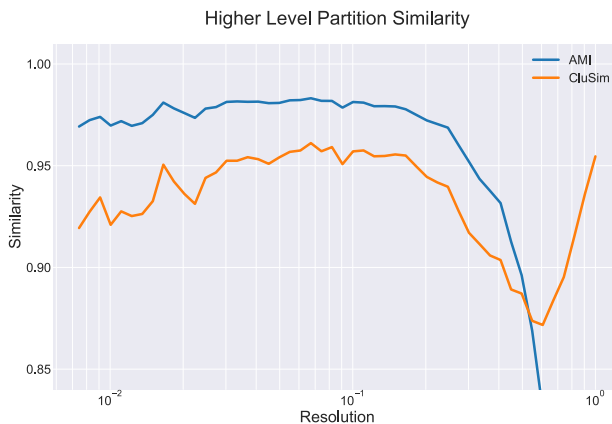


Fig. 8 Higher-level network partition stability with varying resolution. Partitions at smaller resolutions than those shown in the figure tend towards a single community and partitions at larger resolutions tend towards a unique community for each node. There are maxima for each similarity measure around $r = 0.067$

Appendix D: Topic labelling

Figure 9 shows the interface used by the independent labellers.

```

Cluster 9
Most Frequently Occurring Core Articles:
European Union      49
European Commission 12
European Parliament  5
Donald Tusk         3
European Council    2
dtype: int64

Most Frequently Occurring Articles in Networks:
European Union      59
European Commission 22
European Parliament 22
European Council    20
Council of the European Union 18
dtype: int64

Random Sample of Events:

20180212
Cyprus–Turkey maritime zones dispute
The European Union urges calm and restraint after Turkish Navy warships obstruct a Cypriot offshore
drilling vessel in the Eastern Mediterranean, which was approaching an area to explore for natural
gas. (Reuters)

20181121
Kosovo puts a 100% trade tariff on all goods imported from Serbia and Bosnia and Herzegovina. The
European Union says this is a "clear violation" of the Central European Free Trade Agreement. (BBC)

20180824
European migrant crisis
Representatives from 12 European Union countries do not reach an agreement on Italy's request to take
in the 150 migrants still remaining on the Diciotti-class vessel in Catania, Sicily, after their
rescue by the Italian Coast Guard nine days ago. (Deutsche Welle) (ANSA)

20180904
The European Ombudsman, Emily O'Reilly, mentions four counts of maladministration by the European
Commission in the fast-track nomination of Martin Selmayr as its Secretary-General in February. (BBC)

20181027
ALDE European Parliament group expels Catalan nationalist party PDECat amid corruption scandals that
affect the CDC predecessor party. (El Periódico)

Would you like to see more event descriptions? (Y/N):
n

Describe this cluster:

```

Fig. 9 The interface for labelling Topics of Attention, showing the most frequently occurring core articles, regular articles, and a sample of related events

Data availability statement The code used and data generated during the current study are available in the WikiNewsTopics GitHub repository: <https://github.com/pgilders/WikiNewsNetwork-01-WikiNewsTopics>.

Declarations

Conflict of interest The authors declare no conflicts of interest concerning this piece of research.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is

not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Rosengren, K. E. (1970). International news: Intra and extra media data. *Acta Sociologica*, *13*(2), 96–109.
2. Moat, H. S., Curme, C., Avakian, A., Kenett, D. Y., Stanley, H. E., & Preis, T. (2013). Quantifying Wikipedia usage patterns before stock market moves. *Scientific Reports*, *3*(1), 1–5.
3. Mestyán, M., Yasseri, T., & Kertész, J. (2013). Early prediction of movie box office success based on Wikipedia activity big data. *PLoS ONE*, *8*(8), 71226.
4. McIver, D. J., & Brownstein, J. S. (2014). Wikipedia usage estimates prevalence of influenza-like illness in the United States in near real-time. *PLoS Computational Biology*, *10*(4), 1003581.
5. Georgescu, M., Pham, D.D., Kanhabua, N., Zerr, S., Siersdorfer, S., & Nejdil, W. (2013). Temporal summarization of event-related updates in Wikipedia. In: *Proceedings of the 22nd International Conference on World Wide Web*, pp. 281–284.
6. Ahn, B.G., Van Durme, B., & Callison-Burch, C. (2011). WikiTopics: What is popular on Wikipedia and why. In: *Proceedings of the Workshop on Automatic Summarization for Different Genres, Media, and Languages*, pp. 33–40. Association for Computational Linguistics.
7. Keegan, B.C., Gergle, D., & Contractor, N. (2011). Hot off the wiki: Dynamics, practices, and structures in Wikipedia's coverage of the tōhoku catastrophes. In: *Proceedings of the 7th International Symposium on Wikis and Open Collaboration*, pp. 105–113.
8. Kobayashi, R., Gildersleve, P., Uno, T., & Lambiotte, R. (2021). Modeling collective anticipation and response on Wikipedia. In: *Proceedings of the International AAAI Conference on Web and Social Media*, vol. 15, pp. 315–326.
9. Miz, V., Ricaud, B., Benzi, K., & Vanderghenst, P. (2019). Anomaly detection in the dynamics of web and social networks using associative memory. In: *The World Wide Web Conference*, pp. 1290–1299.
10. Miz, V., Hanna, J., Aspert, N., Ricaud, B., & Vanderghenst, P. (2020). What is trending on wikipedia? capturing trends and language biases across wikipedia editions. *Companion Proceedings of the Web Conference, 2020*, 794–801.
11. Keegan, B.C. (2012). High tempo knowledge collaboration in Wikipedia's coverage of breaking news events. PhD Thesis.
12. Gildersleve, P. (2021). The wikipedia news network: understanding collective response to current events through the internet's encyclopaedia. PhD thesis, University of Oxford.
13. Rosenzweig, R. (2006). Can history be open source? Wikipedia and the future of the past. *The Journal of American History*, *93*(1), 117–146.
14. Pentzold, C. (2009). Fixing the floating gap: The online encyclopaedia Wikipedia as a global memory place. *Memory Studies*, *2*(2), 255–272.
15. Ferron, M., Massa, P.: Collective memory building in Wikipedia: The case of North African uprisings. In: *Proceedings of the 7th International Symposium on Wikis and Open Collaboration*, pp. 114–123.
16. Kanhabua, N., Nguyen, T.N., & Niederée, C. (2011). What triggers human remembering of events?: A large-scale analysis of catalysts for collective memory in Wikipedia. In: *Proceedings of the 14th ACM/IEEE-CS Joint Conference on Digital Libraries*, pp. 341–350 (2014). IEEE Press.
17. Luyt, B. (2015). Wikipedia, collective memory, and the Vietnam War. *Journal of the Association for Information Science and Technology*.
18. García-Gavilanes, R., Mollgaard, A., Tsvetkova, M., & Yasseri, T. (2017). The memory remains: Understanding collective memory in the digital age. *Science Advances*, *3*(4), 1602368.
19. Twyman, M., Keegan, B.C., & Shaw, A. (2017). Black Lives Matter in Wikipedia: Collective memory and collaboration around online social movements. In: *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*, pp. 1400–1412. ACM.

20. Candia, C., Jara-Figueroa, C., Rodriguez-Sickert, C., Barabási, A.-L., & Hidalgo, C. A. (2019). The universal decay of collective memory and attention. *Nature Human Behaviour*, 3(1), 82.
21. Yasseri, T., Gildersleve, P., & David, L. (2022). Collective memory in the digital age. *Progress in Brain Research*, 274(1), 203–226.
22. Croatian Wikipedia Disinformation Assessment-2021 (2021). https://meta.wikimedia.org/wiki/Croatian_Wikipedia_Disinformation_Assessment-2021 Accessed 2023-05-08.
23. Grabowski, J., & Klein, S. (2023). Wikipedia's intentional distortion of the history of the holocaust. *The Journal of Holocaust Research*, 1–58.
24. Singer, P., Lemmerich, F., West, R., Zia, L., Wulczyn, E., Strohmaier, M., & Leskovec, J. (2017). Why we read Wikipedia. In: *Proceedings of the 26th International Conference on World Wide Web*, pp. 1591–1600. International World Wide Web Conferences Steering Committee.
25. Giles, J. (2005). Internet encyclopaedias go head to head. Nature Publishing Group. <https://www.nature.com/articles/438900a> Accessed 2023-05-08.
26. Devgan, L., Powe, N., Blakey, B., & Makary, M. (2007). Wiki-surgery? internal validity of Wikipedia as a medical and surgical reference. *Journal of the American College of Surgeons*, 205(3), 76–77.
27. Fallis, D. (2008). Toward an epistemology of Wikipedia. *Journal of the American Society for Information Science and Technology*, 59(10), 1662–1674.
28. Messner, M., & South, J. (2011). Legitimizing Wikipedia: How US national newspapers frame and use the online encyclopedia in their coverage. *Journalism Practice*, 5(2), 145–160.
29. Messner, M., & DiStaso, M. W. (2013). Wikipedia versus Encyclopedia Britannica: A longitudinal analysis to identify the impact of social media on the standards of knowledge. *Mass Communication and Society*, 16(4), 465–486.
30. Thompson, N., & Hanley, D. (2018). Science is shaped by Wikipedia: Evidence from a randomized control trial. Preprint. 10.2139/ssrn.3039505.
31. Yoshida, M., Arase, Y., Tsunoda, T., & Yamamoto, M. (2015). Wikipedia page view reflects web search trend. In: *Proceedings of the ACM Web Science Conference*, pp. 1–2.
32. ElBahrawy, A., Alessandretti, L., & Baronchelli, A. (2019). Wikipedia and cryptocurrencies: interplay between collective attention and market performance. *Frontiers in Blockchain*, 2, 12.
33. Hinnoosaar, M., Hinnoosaar, T., Kummer, M., & Slivko, O. (2021). Wikipedia matters. *Journal of Economics & Management Strategy*.
34. Yasseri, T., & Bright, J. (2014). Can electoral popularity be predicted using socially generated big data? *it-Information Technology* 56(5), 246–253.
35. Yasseri, T., & Bright, J. (2016). Wikipedia traffic data and electoral prediction: Towards theoretically informed models. *EPJ Data Science*, 5(1), 1–15.
36. Matsakis, L. (2018). YouTube Will Link Directly to Wikipedia to Fight Conspiracy Theories. *Conde Nast* <https://www.wired.com/story/youtube-will-link-directly-to-wikipedia-to-fight-conspiracies/> Accessed 2023-05-08.
37. Withers, R. (2018). Amazon Owes Wikipedia Big-Time. *Slate*. <https://slate.com/technology/2018/10/amazon-echo-wikipedia-wikimedia-donation.html> Accessed 2023-05-08.
38. Perez, S. (2020). Facebook tests Wikipedia-powered information panels, similar to Google, in its search results. *TechCrunch*. <https://techcrunch.com/2020/06/11/facebook-tests-wikipedia-powered-information-panels-similar-to-google-in-its-search-results> Accessed 2023-05-08.
39. TwitterInc. (2020). Our plans to relaunch verification and what's next. https://blog.twitter.com/en_us/topics/company/2020/our-plans-to-relaunch-verification-and-whats-next.html Accessed 2023-05-08.
40. Vincent, N., & Hecht, B. (2021). A deeper investigation of the importance of Wikipedia links to search engine results. *Proceedings of the ACM on Human-Computer Interaction* 5(CSCW1), 1–15.
41. Thoppilan, R., De Freitas, D., Hall, J., Shazeer, N., Kulshreshtha, A., Cheng, H.-T., Jin, A., Bos, T., Baker, L., & Du, Y., et al. (2022). Lamda: Language models for dialog applications. arXiv preprint [arXiv:2201.08239](https://arxiv.org/abs/2201.08239).
42. Grimmer, J., & Stewart, B. M. (2013). Text as data: The promise and pitfalls of automatic content analysis methods for political texts. *Political Analysis*, 21(3), 267–297.
43. Boumans, J. W., & Trilling, D. (2016). Taking stock of the toolkit: An overview of relevant automated content analysis approaches and techniques for digital journalism scholars. *Digital Journalism*, 4(1), 8–23.

44. Van Dalen, A., de Vreese, C., & Albæk, E. (2017). Economic news through the magnifying glass: How the media cover economic boom and bust. *Journalism Studies*, 18(7), 890–909.
45. Scharkow, M. (2013). Thematic content analysis using supervised machine learning: An empirical evaluation using German online news. *Quality & Quantity*, 47(2), 761–773.
46. Guo, L., Vargo, C. J., Pan, Z., Ding, W., & Ishwar, P. (2016). Big social data analytics in journalism and mass communication: Comparing dictionary-based text analysis and unsupervised topic modeling. *Journalism & Mass Communication Quarterly*, 93(2), 332–359.
47. Nicholls, T., & Bright, J. (2019). Understanding news story chains using information retrieval and network clustering techniques. *Communication Methods and Measures*, 13(1), 43–59.
48. Aragon, P., Laniado, D., Kaltentbrunner, A., & Volkovich, Y. (2012). Biographical social networks on Wikipedia: A cross-cultural study of links that made history. In: *Proceedings of the Eighth Annual International Symposium on Wikis and Open Collaboration*, pp. 1–4.
49. Kane, G.C. (2009). It's a network, not an encyclopedia: A social network perspective on Wikipedia collaboration. In: *Academy of Management Proceedings*, vol. 2009, pp. 1–6. Academy of Management.
50. Ciampaglia, G. L., Shiralkar, P., Rocha, L. M., Bollen, J., Menczer, F., & Flammini, A. (2015). Computational fact checking from knowledge networks. *PLoS ONE*, 10(6), 0128193.
51. Osborne, M., Petrovic, S., McCreddie, R., Macdonald, C., & Ounis, I. (2012). Bieber no more: First story detection using twitter and wikipedia. In: *Sigir 2012 Workshop on Time-aware Information Access*, pp. 16–76. Citeseer.
52. García-Gavilanes, R., Tsvetkova, M., & Yasseri, T. (2016). Dynamics and biases of online attention: The case of aircraft crashes. *Royal Society Open Science* 3(10).
53. Kummer, M.E. (2014). Spillovers in networks of user generated content: Pseudo-experimental evidence on Wikipedia. ZEW-Centre for European Economic Research Discussion Paper (14-132).
54. Kummer, M.E. (2013). Spillovers in networks of user generated content—evidence from 23 natural experiments on wikipedia. ZEW-Centre for European Economic Research Discussion Paper (13-098).
55. Zhu, K., Walker, D., & Muchnik, L. (2020). Content growth and attention contagion in information networks: Addressing information poverty on wikipedia. *Information Systems Research*, 31(2), 491–509.
56. Miz, V., Benzi, K., Ricaud, B., & Vanderghenst, P. (2017). Wikipedia graph mining: Dynamic structure of collective memory. arXiv preprint [arXiv:1710.00398](https://arxiv.org/abs/1710.00398).
57. Steiner, T., Van Hooland, S., & Summers, E. (2013). MJ no more: Using concurrent Wikipedia edit spikes with social network plausibility checks for breaking news detection. In: *Proceedings of the 22nd International Conference on World Wide Web*, pp. 791–794.
58. Keegan, B. C., Gergle, D., & Contractor, N. (2013). Hot off the wiki: Structures and dynamics of Wikipedia's coverage of breaking news events. *American Behavioral Scientist*, 57(5), 595–622.
59. Ni, X., Sun, J.-T., Hu, J., & Chen, Z. (2009). Mining multilingual topics from wikipedia. In: *Proceedings of the 18th International Conference on World Wide Web*, pp. 1155–1156.
60. Ni, X., Sun, J.-T., Hu, J., & Chen, Z. (2011). Cross lingual text classification by mining multilingual topics from wikipedia. In: *Proceedings of the Fourth ACM International Conference on Web Search and Data Mining*, pp. 375–384.
61. Syed, Z., Finin, T., & Joshi, A. (2008). Wikipedia as an ontology for describing documents. UMBC Student Collection.
62. Kittur, A., Chi, E.H., & Suh, B. (2009). What's in wikipedia? mapping topics and conflict using socially annotated category structure. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 1509–1512.
63. Boldi, P., & Monti, C. (2016). Cleansing wikipedia categories using centrality. In: *Proceedings of the 25th International Conference Companion on World Wide Web*, pp. 969–974.
64. Johnson, I., Gerlach, M., & Sáez-Trumper, D. (2021). Language-agnostic topic classification for wikipedia. *Companion Proceedings of the Web Conference, 2021*, 594–601.
65. Lehmann, J., Gonçalves, B., Ramasco, J.J., & Cattuto, C. (2012). Dynamical classes of collective attention in Twitter. In: *Proceedings of the 21st International Conference on World Wide Web*, pp. 251–260. ACM.
66. Portal:Current events - Wikipedia (2023). https://en.wikipedia.org/wiki/Portal:Current_events Accessed 2023-05-08.

67. Analytics Datasets: Clickstream (2023). <https://dumps.wikimedia.org/other/clickstream/readme.html> Accessed 2023-05-08.
68. Wikimedia: Wikimedia Downloads (2023). <https://dumps.wikimedia.org/> Accessed 2023-05-08.
69. Hill, B.M., & Shaw, A. (2014). Consider the redirect: A missing dimension of Wikipedia research. In: *Proceedings of The International Symposium on Open Collaboration*, p. 28. ACM.
70. MediaWiki: API:Main page — MediaWiki, The Free Wiki Engine (2023). https://www.mediawiki.org/w/index.php?title=API:Main_page&oldid=3463462 Accessed 2023-05-08.
71. Hagmann, P., Cammoun, L., Gigandet, X., Meuli, R., Honey, C. J., Wedeen, V. J., & Sporns, O. (2008). Mapping the structural core of human cerebral cortex. *PLoS Biology*, 6(7), 159.
72. Michel, C. M., & Koenig, T. (2018). EEG microstates as a tool for studying the temporal dynamics of whole-brain neuronal networks: a review. *NeuroImage*, 180, 577–593.
73. Traag, V.A., Waltman, L., & van Eck, N.J. (2019). From Louvain to Leiden: Guaranteeing well-connected communities. *Scientific Reports* 9.
74. Blondel, V. D., Guillaume, J.-L., Lambiotte, R., & Lefebvre, E. (2008). Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2008(10), 10008.
75. Traag, V. A., Van Dooren, P., & Nesterov, Y. (2011). Narrow scope for resolution-limit-free community detection. *Physical Review E*, 84(1), 016114.
76. Mucha, P. J., Richardson, T., Macon, K., Porter, M. A., & Onnela, J.-P. (2010). Community structure in time-dependent, multiscale, and multiplex networks. *Science*, 328(5980), 876–878.
77. Page, L., Brin, S., Motwani, R., & Winograd, T. (1999). The PageRank citation ranking: Bringing order to the web. Technical report, Stanford InfoLab. <http://ilpubs.stanford.edu:8090/422/>.
78. Ioffe, S. (2010). Improved consistent sampling, weighted minhash and l1 sketching. In: *2010 IEEE International Conference on Data Mining*, pp. 246–255. IEEE.
79. Galtung, J., & Ruge, M. H. (1965). The structure of foreign news: The presentation of the Congo, Cuba and Cyprus crises in four Norwegian newspapers. *Journal of Peace Research*, 2(1), 64–90.
80. Shoemaker, P. J., Chang, T.-K., & Brendlinger, N. (1987). Deviance as a predictor of newsworthiness: Coverage of international events in the US media. *Annals of the International Communication Association*, 10(1), 348–365.
81. Harcup, T., & O’Neill, D. (2001). What is news? *Galtung and Ruge revisited*. *Journalism Studies*, 2(2), 261–280.
82. Harcup, T., & O’Neill, D. (2017). What is news? News values revisited (again). *Journalism Studies*, 18(12), 1470–1488.
83. Graham, M., Hogan, B., Straumann, R. K., & Medhat, A. (2014). Uneven geographies of user-generated information: Patterns of increasing informational poverty. *Annals of the Association of American Geographers*, 104(4), 746–764.
84. Callahan, E. S., & Herring, S. C. (2011). Cultural bias in Wikipedia content on famous persons. *Journal of the American Society for Information Science and Technology*, 62(10), 1899–1915.
85. Hecht, B., & Gergle, D. (2009). Measuring self-focus bias in community-maintained knowledge repositories. In: *Proceedings of the Fourth International Conference on Communities and Technologies*, pp. 11–20. ACM.
86. Adams, J., Brückner, H., & Naslund, C. (2019). Who counts as a notable sociologist on Wikipedia? gender, race, and the “professor test”. *Socius*, 5, 2378023118823946.
87. Wikipedia Founder Jimmy Wales Responds. Slashdot.org (2004). <https://slashdot.org/story/04/07/28/1351230/wikipedia-founder-jimmy-wales-responds> Accessed 2023-05-08.
88. Boydston, A. E., Hardy, A., & Walgrave, S. (2014). Two faces of media attention: Media storm versus non-storm coverage. *Political Communication*, 31(4), 509–531.
89. Bao, P., Hecht, B., Carton, S., Quaderi, M., Horn, M., & Gergle, D. (2012). Omnipedia: Bridging the Wikipedia language gap. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 1075–1084.
90. Hale, S.A. (2014). Multilinguals and Wikipedia editing. In: *Proceedings of the 2014 ACM Conference on Web Science*, pp. 99–108. ACM.
91. Lemmerich, F., Sáez-Trumper, D., West, R., & Zia, L. (2019). Why the world reads Wikipedia: Beyond English speakers. In: *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*, pp. 618–626. ACM.
92. Welsler, H.T., Cosley, D., Kossinets, G., Lin, A., Dokshin, F., Gay, G., & Smith, M. (2011). Finding social roles in Wikipedia. In: *Proceedings of the 2011 iConference*, pp. 122–129. ACM.

93. West, R., Weber, I., & Castillo, C. (2012). A data-driven sketch of Wikipedia editors. In: *Proceedings of the 21st International Conference on World Wide Web*, pp. 631–632. ACM.
94. Yang, D., Halfaker, A., Kraut, R., & Hovy, E. (2016). Who did what: Editor role identification in Wikipedia. In: *Tenth International AAAI Conference on Web and Social Media*.
95. Wikipedia: How the Current events page works (2023). https://en.wikipedia.org/wiki/Wikipedia:How_the_Current_events_page_works Accessed 2023-05-08.
96. Milne, D., & Witten, I.H. (2008). Learning to link with Wikipedia. In: *Proceedings of the 17th ACM Conference on Information and Knowledge Management*, pp. 509–518.
97. Wikipedia: In the news (2023). https://en.wikipedia.org/wiki/Wikipedia:In_the_news Accessed 2023-05-08.
98. Wikipedia: Redirect (2023). <https://en.wikipedia.org/wiki/Wikipedia:Redirect> Accessed 2023-05-08.
99. Analytics/Data_Lake/Traffic/Pageviews/Redirects (2023). https://wikitech.wikimedia.org/wiki/Analytics/Data_Lake/Traffic/Pageviews/Redirects Accessed 2023-05-08.
100. Gates, A. J., & Ahn, Y.-Y. (2019). CluSim: a python package for calculating clustering similarity. *Journal of Open Source Software*, 4(35), 1264.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Authors and Affiliations

Patrick Gildersleve¹  · Renaud Lambiotte²  · Taha Yasseri³ 

✉ Patrick Gildersleve
p.gildersleve@lse.ac.uk
<https://scholar.google.co.uk/citations?user=3CTTLDgAAAAJ>
Renaud Lambiotte
<https://scholar.google.co.uk/citations?user=IftXi28AAAAJ>
Taha Yasseri
<https://scholar.google.co.uk/citations?user=AiLoMKwAAAAJ>

¹ Department of Methodology, London School of Economics and Political Science, Houghton Street, London WC2A 2AE, UK

² Mathematical Institute, University of Oxford, Oxford OX2 6GG, UK

³ School of Sociology, University College Dublin, Dublin 4, Ireland