

Cross-Language Evolution of Divergent Collective Memory Around the Arab Spring

H. Laurie Lawson, Brian C. Keegan

University of Colorado Boulder

Laurie.lawson@colorado.edu, Brian.Keegan@colorado.edu

Abstract

The Arab Spring was a cascade of protests and uprisings across the Middle East and North Africa that toppled governments and escalated into prolonged conflicts in others. Collective memories of these events evolve unevenly across political, cultural, and linguistic contexts. Drawing on archived content from the Arabic and English Wikipedias (2011–2024), this study provides a generalizable framework for studying long-term collective memory processes on Wikipedia. We propose multilingual measures of salience, deliberation, contextualization, and consolidation and evaluate them using English and Arabic Wikipedia articles about the Arab Spring. Article size and wikilinks indicate sustained salience in both languages, though along different trajectories. Temporal clustering of wikilink inclusion reveals *Stable*, *Debated*, and *Forgotten* content, marking zones of consensus, contestation, and exclusion. Interlanguage links show limited cross-lingual alignment, with only a minority shared between editions. English pages consistently link back to the Arab Spring summary, while Arabic pages often do not, yielding uneven integration of the phenomenon into the encyclopedic network. These findings demonstrate Wikipedia’s role in long-term memory work and highlight cross-lingual biases relevant to multilingual knowledge systems.

Code — <https://github.com/hlaurielawson/Cross-Language-Evolution-of-Divergent-Collective-Memory-Around-the-Arab-Spring/tree/main>

Background

Collective memory refers to the socially produced ways groups construct and negotiate shared understandings of the past (Halbwachs 1992). Memory scholars emphasize that such understandings are not static records but are continually shaped through processes of selection, omission, reinterpretation, and stabilization over time (Assmann and Czaplicka 1995; Zubrzycki and Woźny 2020). These processes are mediated by institutions, technologies, and cultural practices that determine what becomes durable, visible, or forgotten (Nora 1989; Kansteiner 2002). Empirically tracing collective memory is therefore challenging, particularly across long time horizons and multiple linguistic contexts.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Digital platforms have become central sites of contemporary memory production and analysis by providing infrastructures that both mediate and archive memory work (Hoskins 2017). Wikipedia, in particular, has emerged as a prominent test case for studying collective memory and is often described as a “global memory place” (Pentzold 2009; Yasseri, Gildersleve, and David 2022). Unlike traditional reference works, Wikipedia does not require different language editions to present identical or even similar accounts of the same topic. This structural openness has produced distinct editorial norms, practices, and resource environments across language editions (Bipat, McDonald, and Zachry 2018; Hara, Shachaf, and Hew 2010; Nemoto and Gloor 2011). As a result, Wikipedia does not reflect a unified global memory or the memories of entire linguistic or national community, but rather the situated collective memory work of editor communities and automated agents shaped by language, access to sources, platform affordances, and governance regimes (Roy, Bhatia, and Jain 2022; Ferron and Massa 2011a; Porter, Krafft, and Keegan 2020). Within this framing, collective memory is not just the content that is produced, but the result of sociotechnical infrastructures and institutional processes. This makes Wikipedia a critical site for examining how collective memory evolves unevenly across linguistic and cultural contexts, directly supporting our goal of tracing long-term memory construction across languages.

Because every edit is preserved in publicly accessible revision histories, Wikipedia offers a rare empirical record of collective memory in practice, enabling longitudinal and cross-linguistic analyses of how narratives are formed, contested, and stabilized over time (Ford 2022). Prior research has leveraged this infrastructure to examine article structure (Porter, Krafft, and Keegan 2020; Hickman et al. 2021; He et al. 2018; Dandala, Mihalcea, and Bunescu 2012), content (Hecht and Gergle 2010; Roy, Bhatia, and Jain 2022; Massa and Scrinzi 2012), and collaboration (Hale 2014). Studies of historical events on Wikipedia, however, tend to emphasize immediate aftermaths, such as bursts of editing or attention, rather than the longer-term evolution of collective memory (Ferron and Massa 2011a; Luyt 2016; Twyman, Keegan, and Shaw 2017). More recent computational approaches model short-term event clustering (Gildersleve, Lambiotte, and Yasseri 2023), attention decay (Igarashi et al. 2022), or structural inequities in coverage (Rupprechter,

Burghardt, and Helic 2023), but still rarely capture how events are integrated into broader historical narratives over time.

This study addresses that gap by introducing a generalizable analytic framework for examining the afterlife of Wikipedia articles about historic events, using the Arab Spring as a case to trace how collective memory unfolds years after the initial upheaval. We focus on the longitudinal evolution of two types of links within the English and Arabic articles on the Arab Spring. The first are wikilinks, which connect articles and signal editorial judgments of relevance and contextual association¹. The second are inter-language links (ILLs), which identify corresponding topics across language editions². By analyzing these link structures over time, we compare how concepts are introduced, removed, and made available across languages, using link dynamics as proxies for collective memory processes.

While links do not fully summarize article content, editorial consensus to create and maintain them provides a strong structural signal of how concepts are related and situated over time. This longitudinal structural perspective complements content analyses of article text and talk page discussions but also generalizes across language editions, enabling comparisons of how events are linked, situated, and re-defined within Wikipedia's multilingual ecosystem. This structural focus is particularly valuable given the evolving and uneven quality of large-scale content analysis methods for non-Latin scripts such as Arabic. Building on this approach, our framework presents an ensemble of quantitative methods to develop a grounded understanding of four behaviors related to collective memory formation: *salience*, *deliberation*, *contextualization*, and *consolidation*.

Our study is guided by four research questions, each linked to a methodological strategy:

1. **Salience:** How does sustained editorial investment in the Arab Spring's narrative, manifest over time and across language editions on Wikipedia?

We track changes in article size and the number of wikilinks as indicators of editorial engagement or neglect.

2. **Deliberation:** How do patterns of editorial deliberation across languages shape divergent perspectives on the Arab Spring?

We apply temporal clustering of wikilink inclusion to identify patterns of *Stable*, *Debated*, and *Forgotten* content.

3. **Contextualization:** How do the English and Arabic articles differ in the concepts they use to contextualize the Arab Spring within broader historical and political narratives?

We analyze the inter-language links (ILLs) of wikilinks to detect overlaps and asymmetries in the concepts available and actually referenced across editions.

4. **Consolidation:** How and when do related articles across Wikipedia link back to the Arab Spring, integrating it into a wider encyclopedic memory of the period?

¹https://en.wikipedia.org/wiki/Wikipedia:Manual_of_Style/Linking

²https://en.wikipedia.org/wiki/Help:Interlanguage_links

We generate ego-networks of related articles to capture the degree and timing of the Arab Spring's integration into wider Wikipedia coverage.

Longitudinal Measures of Event Salience

To examine how sustained editorial investment or salience, in the Arab Spring's narrative manifests over time and across language editions on Wikipedia, we analyze longitudinal changes in article size and the number of wikilinks. Article size reflects the amount of content devoted to the event, while wikilinks capture how editors situate it within broader historical and political contexts. Consistent with theories of collective memory that emphasize durability and stabilization alongside moments of reinterpretation, sustained editorial investment may take the form of periods of growth as well as prolonged stability, reflecting the maintenance of a consolidated narrative rather than ongoing revision (Halbwachs 1992; Assmann and Czaplicka 1995). For this study, we define an article as "neglected" if it experiences no substantial content or structural revisions for more than two consecutive years, aside from routine maintenance. We get this data through the Wikipedia API. Figure 1 presents the changes in article size and wikilink counts as our operationalization of salience.

Size

In our test case, article size captures how much content editors devote to the Arab Spring over time. According to (Roy, Bhatia, and Jain 2020) about 65% of articles are longer in English than in Arabic, meaning the overall discrepancy in size between the two articles is not unexpected.

As seen in Figure 1a, size fluctuates sharply in the English article, reflecting impactful restructurings. The first major decline occurred in 2012–2013, when the page was converted into a summary article and content was moved to a concurrent incidents page. A second dip in 2016 stemmed from large reference edits and restructuring of the *Social Media and the Arab Spring* subsection, followed by a rebound as editors expanded both the *Aftermath* and *Social Media* sections. In early 2017, a bot intervention produced a sharp increase by repairing nearly 200 sources, while in 2019 a single editor restructured most of the page, adding or altering five sub sections of the article.

By contrast, the Arabic article shows fewer large disruptions. After a noticeable decline in 2013, when interwiki links³ were deleted, the only other major disruption occurred in 2018, when a bot repaired broken external references. Beyond these moments, growth has been gradual but consistent, with small editorial changes continuing into 2023, such as revisions to politically sensitive phrasing about the Syrian conflict. Taken together, Figure 1a shows that salience is maintained in English through dramatic restructurings, while in Arabic it is sustained by steady, incremental adjustments.

³Interwiki links are a link to another project, to another language and both, to another project in another language. This is compared to a wikilink that only links within the same wiki.

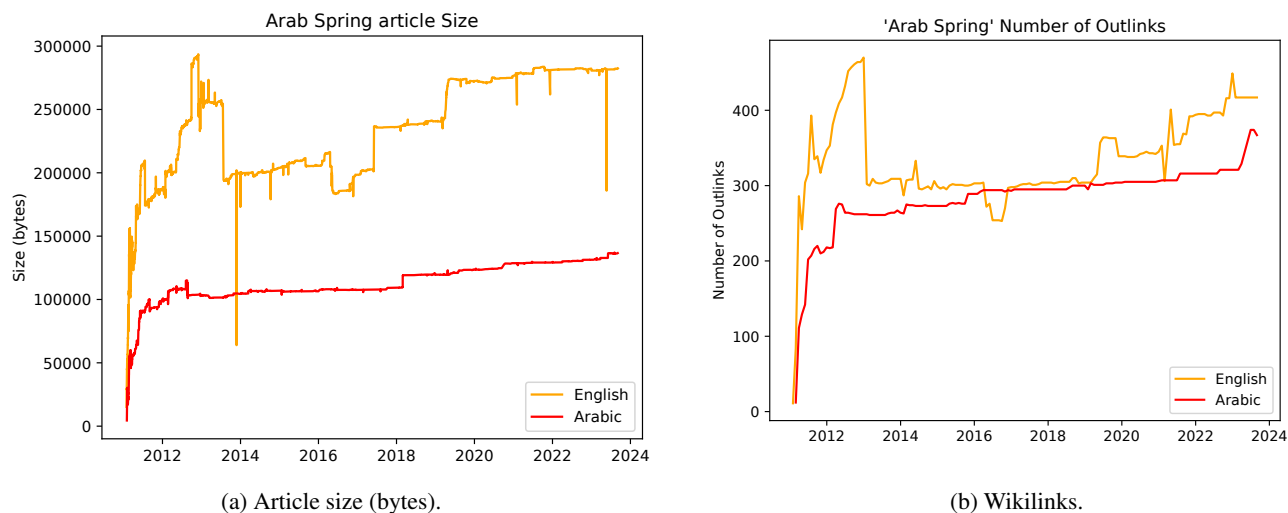


Figure 1: Changes in the size (left) and number of wikilinks (right) for the Arab Spring article in English and Arabic used to operationalize salience of the article through time.

Number of Wikilinks

The number of wikilinks captures other types of engagement that is related to the collective memory processes such as the linking of related concepts within broader contexts (Roy, Bhatia, and Jain 2022). Since the creation of Wikipedia, the number of wikilinks per article has increased over time, partly due to the templates used to structure the article but also due to editing norms (Consonni, Laniado, and Montresor 2019; Buriol et al. 2006). The wikilinks that are added within templates adds a compounding variable; however, we mitigated this by only including the wikilinks that are within the text of the article.

In the English article, the overall trajectory predictably mirrors changes in size: sharp disruptions tied to large editorial interventions. Outside the largest drop in 2013, when wikilinks were moved during the shift to a summary article, the most notable changes occurred in 2016 and 2019. In 2016, two editors restructured the *Social Media and the Arab Spring*, *International Pressures*, and *Aftermath* subsections, producing an inverted plateau. In 2019, a single editor revised six subsections including *Results*, *Aftermath*, *Social Justice*, and *Counter-revolution and Civil Wars*, leading to a sharp increase. Smaller attempts to revise the summary box in 2021 and to add 23 links in 2022 were quickly undone, underscoring the contested nature of these changes.

In the Arabic article, wikilink variation is more modest but still present. The largest increase occurred in February 2012, when an editor comprehensively overhauled the introduction. Two smaller increases followed: in 2015 with additions about Jordan, and in 2021 when a bot and editor contributed material on Morocco. The final uptick came in 2023, when a single editor made extensive revisions tagged as *تحرير مرئي* (Visual Editing). Figure 1b shows that, much like size, English salience through wikilinks is maintained through punctuated bursts, while Arabic salience is reflected

in smaller but persistent wikilink additions.

Taken together, these patterns show that the Arab Spring remains a salient topic across both English and Arabic Wikipedia from 2011 through early 2024. In English, salience is evident in repeated large-scale restructurings and bursts of editorial activity, while in Arabic it appears in steadier but persistent small-scale updates. Although the forms differ, neither article has become “settled” or “neglected.” Instead, both remain under active editorial negotiation, confirming that the Arab Spring continues to function as a site of evolving collective memory across linguistic communities.

Temporally Clustered Vectors

Research on collective memory in Wikipedia has largely emphasized deliberation as observable activity, focusing on discussion pages, edit histories, and content change, while paying less attention to how deliberation is ultimately encoded in the platform’s structural features (Ferron and Massa 2011a; Keegan 2019; Hickman et al. 2021; Hale 2014). In this section, we use wikilink inclusion and exclusion as a proxy for deliberation outcomes. We recognize that links may be added or removed for subjective or stylistic reasons that do not directly signal changing interpretations. However, in the aggregate, temporal patterns of wikilink persistence and disappearance reveal which connections editors collectively sustain and which they allow to fade, offering insight into how relevance and emphasis shift over time. Wikilinks are a core feature of Wikipedia’s infrastructure, designed to guide readers toward related concepts editors consider important for contextualizing a topic.

With every revision to a Wikipedia article, Wikipedia creates a new version of the article that includes all of its metadata (such as wikilinks, sources, template structure, etc.) and marks that version of the article with a revision id. We pulled

all of the wikilinks of all of the revision ids of the “Arab Spring” article in English and Arabic and created a binary vector of inclusion for each outlink, also referred to as one-hot encoding. We then computed the pairwise cosine distances between each binary vector within each language and clustered through hierarchical clustering due to its flexibility and readability. We evaluated the fidelity of the hierarchical clustering using the cophenetic correlation coefficient (CCC), which measures how well the dendrogram preserves the original pairwise distance relationships. The resulting CCC of 0.69 indicates that the hierarchical structure provides a faithful representation of the underlying similarity space. We then qualitatively analyzed the temporally clustered wikilinks to identify emergent themes, narrative emphases, and patterns of inclusion that reflect evolving collective memory structures.

Figures 2a and 3a present pairwise similarity heatmaps, which group revisions based on the inclusion patterns of the wikilinks through time. The rows and columns are the superset of wikilinks across all versions of an article and the values are cosine similarities. This clustering surfaces distinct trajectories of inclusion, exclusion, and persistence that correspond to editorial consensus or contestation over time. Figures 2b and 3b complements this view by plotting the number of links within each cluster across the full timeline, showing when clusters rise, decline, or disappear. Taken together, the heatmap identifies which groups of wikilinks cohere as meaningful clusters, while the line chart reveals their temporal dynamics, allowing us to distinguish *Stable*, *Debated*, and *Forgotten* concepts. This combined approach makes visible both the structural similarity of revisions and the longitudinal patterns through which collective memory is negotiated.

English Deliberation Around Wikilink Inclusion

Cluster 1 This cluster has identified what we label to be *Forgotten* concepts. By 2020, all of the wikilinks in this cluster were removed and are never added again. In 2016, an editor came in and deleted the sections on international reactions and external effect but added the section *Disparate Outcomes*. This resulted in links to articles about smaller events such as “2011–12 Moroccan protests” were completely forgotten. A second contraction occurred in late 2019, when revisions to *The Aftermath* and country-specific sections (Yemen, Libya, Syria) further reduced regional context. Taken together, these deletions illustrate how certain international and regional perspectives became forgotten, narrowing the scope of the Arab Spring article over time.

Cluster 2 We designate cluster 2 as *Stable* because the articles in this group rarely undergo deletions and constitute the majority of current content. Initially anchored by core wikilinks such as the “2011 Egyptian Revolution” and “2011 Syrian revolution”, this cluster has remained relatively consistent in scope while gradually expanding its orientation toward Western perspectives. This shift is evident in the integration of wikilinks connected to political philosophy, including “Foreign policy,” “Media Activism,” and “Barack Obama”, which signal a broadening of the narrative contextualization beyond the regional events themselves.

Cluster 3 Cluster 3, which we term *Debated*, is defined by recurrent cycles of contested edits in which newly introduced links are quickly removed, reflecting ongoing disputes over how the Arab Spring should be contextualized. For instance, in 2014, links to articles such as “Civil Disobedience” and “Rebellion” were briefly inserted into the *Causes* section before being undone. Later spikes followed a similar pattern: in 2022 and 2023, attempts to add entries like “BBC” and “CNN” were swiftly reverted, though in some cases additions such as “Insurgency” and the “Khartoum massacre (2019)” persisted. Plateau periods between 2015 and 2020 reflect longer cycles of addition and removal around wikilinks to articles surround broader regional or transnational events such as “Western Sahara” and “Parliament of India”.

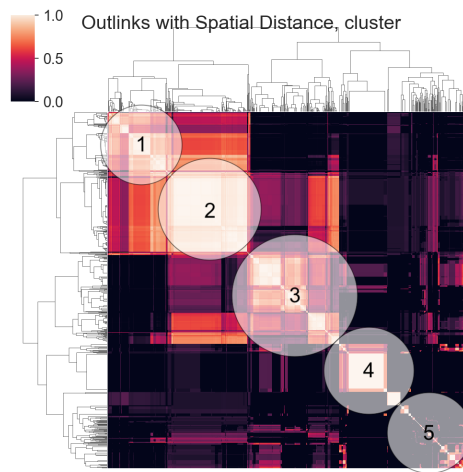
These patterns are driven largely by anonymous edits, often originating from mobile devices, which are typically classified as “vandalism” under Wikipedia’s content policies.⁴ However their frequency points to underlying contestation: while some editors enforce a relatively bounded interpretation of the Arab Spring’s causes, other contributors repeatedly attempt to foreground alternative factors and framings.

Cluster 4 and 5 We analyze clusters 4 and 5 together, as both fall under the category of *Forgotten* concepts. These clusters represent elements that were actively incorporated in the early evolution of the Arab Spring article but subsequently removed, with nearly all traces gone by 2014. Their evolution exemplifies how temporal scope shapes collective memory research. These clusters capture the only content that would appear in previous research, limiting attention to the moment of peak activity.


Cluster 4 is marked by a sharp spike in March 2011, when links related to countries outside the Middle East such as the “Chinese Communist Party”, “Republic of Korea Armed Forces”, and “South Sudan” were briefly added before being deleted by the end of the month. These inclusions reflected discussions of whether the Arab Spring should be understood as part of a broader, global wave of protest. For instance, leaflets referencing the Arab Spring were reportedly dropped in North Korea by the South Korean military during this period. Such material was later relocated to a separate article on the Impacts of the Arab Spring. The removal of these references illustrates early editorial decisions to bound the phenomenon geographically, excluding parallels and ripple effects outside the region.

Cluster 5 follows a similar path but focuses on intra-regional *Forgotten* concepts. While Cluster 4 reflects exclusions beyond the Middle East, Cluster 5 documents the erasure of certain Middle Eastern contexts themselves. By 2012–2014, links to events such as the “2011 Lebanese Protests”, “The Cabinet of Yemen”, and references to Palestine (“Palestinian Intifada”, “Gaza Strip”) or the Gulf (“Royal Oman Police”, “United Arab Emirates”) were removed. Much of this content disappeared when an editor deleted the *Concurrent Incidents* section of this article and put it into a different Wikipedia page. This restructuring nar-

⁴https://en.Wikipedia.org/wiki/Vandalism_on_Wikipedia



(a) This is a pairwise similarity clustered heatmap of wikilink temporal inclusion vectors for revision editions of the English Arab Spring article. This chart includes all wikilinks over time; however, individual wikilink identities are not labeled on the axes due to spatial constraints.



(b) These are the number of wikilinks included through time within each clusters of the English Arab Spring heatmap, the x-axis is normalized by maximum values

Figure 2: English wikilink similarities.

rowed the definition of which countries and events were considered part of the Arab Spring, leaving some details embedded in broader summaries but excising others altogether.

Together, Clusters 4 and 5 highlight a critical editorial process: the drawing of boundaries around what counts as part of the Arab Spring. Whereas Cluster 4 shows the exclusion of global resonances, Cluster 5 reveals a selective forgetting within the region itself, shaping a canonical narrative that privileges certain national experiences while erasing others.

Arabic Deliberation Around Wikilink Inclusion

Cluster 1 Cluster 1 in the Arabic heatmap, analogous to the *Stable* cluster in English, encompasses wikilinks that form the backbone of the contemporary article and steadily expand over time. In addition to foundational references to the Arab Spring itself, this cluster includes a distinctive infrastructure of date-specific links. For example, 14 يناير (January 14), marks major protests in Tunisia, Lebanon, and Jordan but links to a Wikipedia page that includes links to articles related to that day every year.

As the article evolves, the cluster also absorbs references to subsequent regional developments, extending the temporal frame of the Arab Spring. Notable additions include

احتجاجات الجيش السوري الحر (Free Syrian Army) in 2014,


جنوب العراق ٢٠١٨ (Southern Iraqi Protests 2018), and أحزاب سياسية في المغرب (Political Parties in Morocco) in 2021. These inclusions illustrate an ongoing editorial consensus: the Arabic article is not frozen at the moment of the uprisings but is continually updated and contextualized

within subsequent waves of protest and political transitions


Cluster 2 Cluster 2 reveals two distinct phases of inclusion, before 2014 and after 2021, with no overlap between them. For this reason, we classify it as both *Forgotten* and *Debated*. In the early phase, wikilinks tied to neighboring or contemporaneous contexts such as أوغندا (Uganda), إثيوبيا (Ethiopia), الشباب (تنظيم صومالي) (Al Shabaab), and حسن مشيمع (HaSan Mushaima, a Bahraini opposition leader). These “forgotten” entries disappeared in a large edit by an editor who noted that the Arabic version was the only language edition to frame protests in Somalia as part of the Arab Spring, and who deleted the associated section and links.

The second phase emerges after 2021, when wikilinks such as أحزاب سياسية في المغرب (Political Parties in Morocco) and دستور المغرب (Moroccan Constitution) were introduced. These remain in the current version, signaling continued debate over the boundaries of the Arab Spring and the extent to which regional developments should be incorporated into its summary page and thus these articles become labeled as *Debated*. In June of 2021, without a description of the edit in the records of the article, an editor re-writes the descriptions of the events in Morocco, excluding its ties to other protests and focusing more on internal dynamics.

In sum, the clustered patterns of wikilink inclusion and exclusion reveal how Wikipedia’s coverage of the Arab Spring has been continuously shaped by processes of negotiation and boundary-setting. *Stable* clusters capture areas of enduring consensus, where certain events or perspectives become anchored as part of the canonical narrative. *Forgotten* clusters highlight content that was once included but later excluded, reflecting editorial decisions to



(a) This is a pairwise similarity clustered heatmap of wikilink temporal inclusion vectors for revision editions of the Arabic Arab Spring article. The chart includes all wikilinks throughout time, however the axis of the graph presents a spatially arranged subset due to spatial restrictions.



(b) These are the number of wikilinks included through time within each clusters of the Arabic Arab Spring heatmap

Figure 3: Arabic Wikilink similarities.

narrow the phenomenon’s scope and delimit what “counts” as the Arab Spring. *Debated* clusters expose the contested terrain, where repeated additions and removals signal unresolved tensions over interpretation and relevance. Taken together, these dynamics illustrate how temporally clustered editing practices function as mechanisms of collective memory, charting both the consolidation and contestation of historical meaning across language editions.

Contextualization While Acknowledging Linguistic Access

Collective memory is not only sustained through what events are recorded, but also through how those events are contextualized in relation to others. On Wikipedia, contextualization is made visible through wikilinks, which signal editorial judgments about what concepts are relevant for situating the Arab Spring. However, contextualization also depends on which resources are available in the first place. Because not all articles exist in every language, inter-language links (ILLs) provide a critical lens for identifying cross-lingual gaps in contextual resources. According to (Bolikowski 2009) around 45% of articles are isolated linguistically, meaning around 45% of the information on Wikipedia is inaccessible if a user is not a multi-linguist. Specifically within English and Arabic (Roy, Bhatia, and Jain 2020) found the Arabic language edition shares about 67.3% of the same topics as the English language edition. Previous work as established that ILLs are a robust way to map concept alignment (Hecht and Gergle 2010) however, we believe it is more of a metric of how topics can be contextualized within language editions. In this sense, contextualization is shaped both by editorial practice (which links are included) and by structural affordances (which links exist).

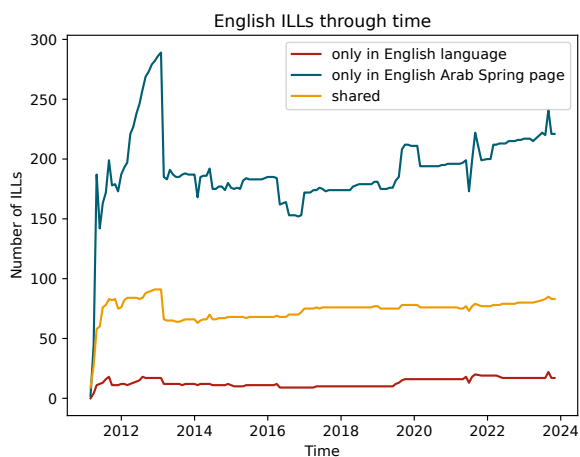
To analyze this process, we leverage the ILLs of wikilinks in the current article to identify discrepancies in what articles are available across English and Arabic Wikipedias. A methodological limitation with it is that Wikipedia’s API does not track the historical evolution of ILLs; we therefore approximate their inclusion based on current availability. One confounding factor that may arise is a redirect⁵⁶ within the ILLs. We resolved all of these variations by just accounting for the link that it navigates to. While this approach cannot capture every temporal nuance, it still allows us to identify systematic differences in the resources available to editors in each language. For cross-language comparability, we group links into three categories. “Only in English language” or “only in Arabic language” refers to articles that exist in just one language. “Only in English Arab Spring page” or “Only in Arabic Arab Spring page” refers to articles that exist in both languages (i.e., have an inter-language link) but are only linked in one version of the Arab Spring article. Finally, “shared” refers to articles that exist in both languages and are linked in both versions of the Arab Spring article.

ILLs of Contemporary Wikilink Contextualization


Table 1 shows that only about one-quarter of English wikilinks (24.2%) and one-third of Arabic wikilinks (29.4%) are shared across languages. A majority of these are links to either major conflicts or people related to the phenomenon such as “2011 Egyptian revolution”, “Bashar al-Assad”, and “Yemeni Revolution”.

⁵<https://en.wikipedia.org/wiki/Wikipedia:Redirect>

⁶Redirects automatically send readers to another page, usually an article or section of an article ex: UK when searched it takes you to “United Kingdom”



(a) English inter-language links.



(b) Arabic inter-language links.

Figure 4: Wikilinks as they evolve through time in English(left) and Arabic(right) categorized by whether their current wikilink has an ILL in the other language version of the article, and whether it is included in that other language version. This is an operationalization of the access and application of shared contextualization practices

	English		Arabic	
	Count	Fraction	Count	Fraction
ILL, wikilinked	79	24.2%	79	29.4%
No ILL	18	5.5%	21	7.8%
ILL, no wikilink	230	70.3%	169	62.8%
midrule Total	327	—	269	—

Table 1: The number of wikilinks by inter-language link (ILL) status for the English and Arabic articles about the Arab Spring.

A majority of the 18 English articles with no Arabic equivalent are American individuals and concepts such as “Zeynep Tufekci”, an American professor studying social media of Arab Spring.

The 21 wikilinks that are only available in Arabic with no correlating article in the English Wikipedia, fall into two main categories: structure-based and content-based. Structure-based wikilinks are links such as, بوابة الشرق (Portal: Middle East)⁷⁸ are articles that are made because of the structure of Wikipedia deemed by the sum of Wikipedia editors in that language edition. Content-based articles such as, محاكمة حسنى مبارك (Trial of Hosni Mubarak) were made because people wanted to have arti-

cles about these topics. This is what all of the 18 articles in the English version with no equivalent in Arabic consist of. Looking at what articles are completely isolated within language editions of Wikipedia, we can see that the English Arab Spring article and the Arabic Arab Spring article reflect clear prioritizations on behalf of the whole Wikipedia linguistic community. However, there are some links that are available in both languages but were chosen not to be referenced in the Arab Spring summary pages. In fact this makes up a majority of the wikilinks.

Isolated in the English version, there continues the pattern of increasing western concepts such as “Condoleezza Rice” as well as concepts about political philosophy such as “Authoritarianism” and “Economic stagnation”. There are also additional wikilinks to information across countries involved in the phenomenon. Such as, “Post-coup unrest in Egypt (2013–2014)”, “Yemeni Crisis (2011–present)”.

Only within the Arabic Wikipedia, there is a large infrastructure of dates. There are links to pages about a singular date but the extensive infrastructure and referencing is only within Arabic. For example, ١٤ يناير (January 14) links to an page that describes everything that happened that day; it is also a major day of protests in Tunisia, Lebanon, and Jordan. Other themes isolated in this section are topics about the Israel-Palestine conflict and wikilinks themed around government structures like such as حرب ١٩٤٨ (1948 War) which English equivalent article is “1948 Arab–Israeli War” and مجلس النواب البحرينى which English article is, “Council of Representatives (Bahrain)”.

As stated when investigating *Deliberation, Stable* wikilinks are what culminate within the current version of the article. Both language articles shared foundational texts. There is a strong inclusion of western concepts and concepts about political philosophy isolated within the English version as well as available across languages but only referenced in En-

⁷<https://en.Wikipedia.org/wiki/Wikipedia:Contents/Portals>

⁸On Wikipedia, a portal “complement main topics in Wikipedia, and expand upon topics by introducing the reader to key articles, images, and categories that further describe the subject and its related topics”

glish. Structural elements within the Arabic language edition of Wikipedia motivate the inclusion of certain wikilinks but there are also many links that focus on Middle Eastern phenomena. Understanding how these links are related to the entire Wikipedia ecosystem, really exemplifies the variation between languages, however to frame this within collective memory processes, we have to take a look temporally at the wikilinks that were included through the consensus building that took place as these articles evolved.

Temporal Wikilink Contextualization

Figure 4(a) shows that the major drop in overall wikilinks observed earlier in Figure 1(b) corresponding to the restructuring of the page into a summary article. This reduction primarily affected links that had inter-language equivalents but were included only in the English version. Because many of these edits centered on the *Social Media* subsection and other “aftermath” sections, the English Arab Spring article came to contextualize post-uprising developments largely through an isolated perspective by drawing on articles available in both languages but referenced exclusively in English.

Seen in Figure 4(b), across time, nearly half of the article’s wikilinks overlap with English at any given time. This high proportion is partly explained by the smaller overall size of the Arabic article, where foundational links account for a larger share of content. As a result, Arabic contextualization remains closely anchored to core concepts, with fewer departures into additional themes. That being said, Figure 4(b) shows that changes in the Arabic article’s wikilinks largely mirror the trends in Figure 1(b). The most significant early shift occurred in 2012, when revisions to the introduction altered many of the shared, foundational links. A smaller spike appears in 2021, when both a bot and an editor added content related to Morocco, much of it absent from the English article, providing an isolated linguistic perspective. In March 2023, another editor made extensive revisions, incorporating links already present in English alongside new ones unique to Arabic.

Taken together, these findings show that contextualization of the Arab Spring diverges sharply across English and Arabic Wikipedias. In English, contextualization has been dynamic but often narrowed by restructurings that foreground Western political concepts and aftermath narratives, even when equivalent articles exist in Arabic. In contrast, the Arabic article anchors more heavily on foundational and regionally specific links, with changes accumulating incrementally rather than through major restructurings. The reliance on date-based infrastructures and Middle Eastern topics in Arabic, versus political philosophy and Western perspectives in English, highlights the asymmetries produced by both editorial choices and structural affordances. These patterns illustrate how linguistic communities construct distinct memory environments, with limited overlap across languages, and set the stage for examining how such differences evolve temporally and shape the collective memory of the *Arab Spring*.

Consolidation through Reference Dissemination

How a community situates a topic in the broader memory system is reflected in how it is referenced when discussing other topics on Wikipedia, influencing how other memories are formed and contextualized. The consolidation of the Arab Spring as a fixture of this system would be its inclusion within the articles about the events and countries involved in the broader phenomenon. We identified eleven events⁹ that both the English and Arabic versions of the summary article deemed “important” and the countries in which those events take place¹⁰. Although prior work shows that pages users perceive as “related” do not always share explicit links, we expect a high degree of link overlap in this case because the Arab Spring is a regional phenomenon that explicitly connects country-specific events (Adafre and de Rijke 2005). As a result, these pages are likely to exhibit a related link structure. More generalizable approaches to relatedness could incorporate additional metrics, such as structural similarity, user-based perspectives, or semantic alignment.

We investigated this consolidation by generating country and event based ego networks¹¹ for monthly page revisions of the articles and their wikilinks and measuring inclusion within those networks and then represented them as gaunt charts to focus on the temporally evolving structure. As editors can link to different articles by hyperlinking any words, not necessarily the name of the article, the resulting figures do not present whether the Arab Spring article is mentioned but whether editors link to it regardless of the words used. This approximates whether editors consider the Arab Spring summary article a necessary and relevant point of information to contextualize the country or event.

Consolidation in English Wikipedia

Figure 5(a) is a visual representation of the inclusion of the Arab Spring article as an wikilink within the English language country articles. The English country articles show relatively little debate, with nearly all including the Arab Spring by 2014. In contrast, the event articles in Figure 5(b) display greater variation, with ongoing debate about whether to reference the summary page. Consolidation is strongest in the Libyan, Bahraini, and Yemeni contexts, where both country and event pages consistently reference the Arab Spring article. These cases however diverge in outcome.

⁹“2011 Omani protests”, “2011–2012 Moroccan protests”, “2011–2012 Jordanian protests”, “2011 Iraqi protests”, “2010–2012 Algerian protests”, “2011 Bahraini uprising”, “Syrian revolution”, “Yemeni Revolution”, “Libyan civil war (2011)”, “2011 Egyptian revolution”, “Tunisian Revolution”

¹⁰These events and countries were cross referenced between languages and the only discrepancy is that in the English article, the event associated with Syria was the “Syrian Revolution” article which is not referenced in the Arabic article at all. Instead in the Arabic article, “The Syrian Civil War” page is referenced. Since that article is also referenced in the English article, that is the event we investigated.

¹¹A type of network from social network analysis that maps connections from a particular perspective




Figure 5: The first graph is a visual of the temporal inclusion of the Arab Spring article as an wikilink in the English language country articles. The second graph is the temporal inclusion of the Arab Spring article in the English language event articles. This presents a perspective on the consolidation of the Arab Spring as a contextualization for the country and event Wikipedia articles.

Bahrain's protests ended by 2018, while in Libya and Yemen the uprisings evolved into prolonged conflicts. Yet all remain closely tied to the Arab Spring narrative. In contrast, Iraqi, Egyptian, and Omani articles show minimal inclusion, suggesting either that editors treated these events as distinct from the Arab Spring, or conversely as so closely tied that additional references were seen as redundant. There is a discrepancy between the *Syria* article including the Arab Spring article consistently from the beginning and the Syrian event not beginning to include or negotiate its inclusion until after 2013. Within English Wikipedia, editors found the Arab Spring a defining characteristic of countries across the MENA¹² region and debate over phenomenon level referencing.

Consolidation in Arabic Wikipedia

The Arabic language country and event articles debate the inclusion of the Arab Spring article from the beginning, as shown in Figure 6. The *Libya* article and *Iraq* article never include the wikilink to the Arab Spring article and the *Egypt* article only includes it for 2 or 3 months at a time; so rarely that it did not show on our visual. Arabic editors did not find linking the Arab Spring article to be a necessary contextualization of these countries. By contrast, articles on countries with smaller uprisings, such as *Jordan* and *Bahrain*, reference the Arab Spring more consistently. *Morocco* and *Oman* show greater fluctuation, incorporating the wikilink only from 2021 onward. The Syrian articles include the Arab Spring article for the longest period overall, though even here debate resurfaces after 2020. Diverging from the English language community on Wikipedia, the Arabic language community found it hardly necessary to link the Arab

Spring article in adjacent articles.

Across both English and Arabic Wikipedias, consolidation patterns reveal how differently linguistic communities situate the Arab Spring article within broader historical narratives. In English, the Arab Spring article is consistently integrated into country articles and frequently referenced in event articles, signaling its role as a central point of contextualization for understanding national and regional developments. By contrast, in Arabic Wikipedia, the Arab Spring article is only intermittently linked to country and event pages, Libya and Iraq never referencing it at all. These disparities suggest that while English editors treat the Arab Spring as a unifying frame across the MENA region, Arabic editors are less inclined to embed the phenomenon article into adjacent narratives, instead allowing events to stand on their own terms. Taken together, these patterns indicate uneven consolidation of the Arab Spring across languages, highlighting how collective memory is not only constructed within articles but also negotiated through the broader web of references that situate events in relation to one another.

Discussion

How do collective memory processes unfold differently across different Wikipedia language editions? Previous research has looked at initial collective memory processes or compared content across language editions, but this paper provides novel empirical findings through and entire temporal timeline and across language. This paper introduces a generalizable analytic framework, operationalized as an ensemble of quantitative constructs for measuring four distinct collective memory processes: salience, deliberation, contextualization, and consolidation. We evaluated these processes using a research design combining variation over time and

¹²Middle East and North Africa




Figure 6: The first graph is a visual of the temporal inclusion of the Arab Spring article as a wikilink in the Arabic language country articles. The second graph is the temporal inclusion of the Arab Spring article in the Arabic language event articles. This presents a perspective on the consolidation of the Arab Spring as a contextualization for the country and event Wikipedia articles.

across languages for articles related to the 2011 Arab Spring.

Salience was operationalized as a measure of the size and number of wikilinks. The English article shows punctuated bursts of restructuring and content expansion, and shows greater variation in size and content over time. While the Arabic article is both shorter and more stable, it evolves gradually with smaller but steady edits.

Deliberation was operationalized as similarity in the temporal vectors of wikilink occurrence. The English article had five distinctive clusters of behavior relating to wikilink inclusion that we classified as *Stable*, *Debated*, and *Forgotten* interpretations while the Arabic article had only two clusters, the first having *Stable* wikilink temporal inclusion and the second cluster having *Forgotten* and *Debated* inclusion patterns. This also revealed that notable changes in size and wikilink inclusion in the Arabic version of the article correlated with *Debated* topics surrounding Morocco and Jordan.

Contextualization was operationalized as a measure of overlap in the inter-language link graph. Despite the availability of relevant topics across languages, the Arabic and English articles referenced very different concepts with English prioritizing Western political concepts and Arabic foregrounding regional conflicts and date infrastructures as seen in their deliberation practices.

Consolidation was operationalized as articles about countries and their national-level protests linking through wikilink to the *Arab Spring* article. Despite the greater geographic and cultural proximity of the events of the Arab Spring to Arabic speakers, the English-language article has far greater dynamism in its content and relationships than the Arabic-language article. The comparative absence of links to الربيع العربي (Arab Spring) article across Arabic-language articles about national histories and even national

events as seen in Figure 6 is particularly surprising.

Our findings directly address a gap in prior work that has largely concentrated on the immediate aftermath of major events rather than their long-term trajectories. As noted in the background, most research on the Arab Spring and on collective memory in Wikipedia has emphasized the days or weeks following protests, when editorial activity is most visible (Twyman, Keegan, and Shaw 2017; Luyt 2016; Ferron and Massa 2011b). By tracing salience, deliberation, contextualization, and consolidation across more than a decade, we extend this lens to reveal how memories continue to be negotiated, forgotten, or stabilized long after the initial surge of attention. This long-term perspective highlights the importance of Wikipedia not just as a record of contemporary events but as a site where historical meaning is continually reworked.

Building on these findings, our comparative approach responds to calls for more systematic cross-lingual analyses of collective memory by showing not only that English and Arabic editions differ in the concepts they emphasize, but also that they sustain or neglect these emphases in diverging temporal patterns. These asymmetries demonstrate how collective memory is shaped by linguistic and cultural contexts, reinforcing the need for multilingual, temporally sensitive methods in digital memory research. Methodologically, our ensemble approach contributes to the growing literature on computational memory studies, which has often emphasized short-term clustering of events (Gildersleve, Lambiotte, and Yasseri 2023), attention decay (Igarashi et al. 2022), or structural inequities in coverage (Rupprechter, Burghardt, and Helic 2023). By integrating temporal clustering of wikilink inclusion, inter-language link (ILL) analysis, and ego-network measures of consolidation, we extend methodological toolkits for examining how memories are stabilized,

contested, or forgotten over more than a decade. This link-based perspective complements text- and discussion-based analyses and offers a replicable framework for future cross-lingual studies of online collective memory.

Implications, Limitations, and Future Work

Our analysis is based on two focal articles, the English and Arabic versions of the Arab Spring, which necessarily limits the scope of our claims. Importantly, our findings reflect the collective memory practices of the editors who contributed to these articles, as expressed through their cumulative editorial decisions over time and is not generalizable across all individuals who speak English or Arabic. While this longitudinal, cross-linguistic design surfaces important differences, we do not claim that these findings can be directly generalized to all cultural representations on Wikipedia. English and Arabic are each global languages with colonial legacies that do not map as cleanly to national boundaries nor does digital affordance utilization map onto direct causation. Instead, they should be read as a case study that illustrates how our framework can be applied and extended to additional language editions and topics and our test case is an example of strong correlation. Moreover, although this corpus reflects the actions of a relatively small and self-selecting group of editors, prior research demonstrates that aggregate editing patterns, link structures, and revision dynamics on Wikipedia serve as durable sociotechnical traces that mediate public knowledge production and collective memory at scale, rendering quantitative metrics not a proxy for individual belief but a method for analyzing how memory becomes institutionally stabilized through platform-mediated practice (Hara, Shachaf, and Hew 2010; Igarashi et al. 2022; Yasseri, Gildersleve, and David 2022; Pentzold 2009)

Another limitation is that our approach primarily analyzes metadata (size, wikilink inclusion/exclusion, and ILLs) rather than the lexical and semantic content of the articles themselves. This allows us to capture structural dynamics in how topics are contextualized, but it may miss important nuances in phrasing, framing, or editorial intention. Similarly, because our design does not incorporate talk page discussions or edit histories, it cannot fully account for the conflicts, negotiations, and subjective decisions that shape article content. In this sense, our results should be viewed as complementary to text-based or discussion-focused approaches rather than as replacements.

Future research should expand this framework to incorporate additional language editions (e.g., Farsi, Hebrew, Turkish) as well as integrate direct text analysis and discussion data. To better capture how memory is constructed and contested in situations. Further, examining whether the “after-lives” of articles documenting major historical events are prone to capture or unreliable narratives as editorial attention wanes would sharpen our understanding of how memory stabilizes or erodes over time. For example, our analysis allows for the inclusion of bots in our generalized analysis which we believe to be a valid inclusion due to the approval neces-

sary to create a bot within wikipedia¹³¹⁴ and the robust infrastructure within both languages for bot monitoring. There are differences in bot usage between language editions but that is outside of the scope of this paper and should be investigated in future work. Finally, because Wikipedia data continue to shape large language models and other downstream systems, identifying and mitigating cross-lingual misalignments is critical for reducing the risk of reproducing biased or fragmented historical narratives in AI.

Taken together, our findings demonstrate that Wikipedia’s treatment of the Arab Spring is not a static archive but an evolving site of memory work, marked by ongoing negotiation, selective forgetting, and uneven consolidation across languages. By tracing salience, deliberation, contextualization, and consolidation over more than a decade, we show how collective memory processes manifest differently in English and Arabic, revealing asymmetries shaped by cultural and linguistic contexts. These results not only extend prior work that has focused on short-term editorial dynamics but also respond to gaps in the literature on long-term, cross-lingual memory practices. Methodologically, our ensemble of link-based measures contributes a replicable framework for studying how historical events are situated, contested, or stabilized in multilingual digital platforms. Substantively, this work underscores the importance of incorporating temporal and cross-linguistic variation into studies of collective memory, both for understanding how societies remember and for evaluating the knowledge infrastructures that inform AI and other downstream systems.

Acknowledgments

This work was partially supported through the Wikimedia Research Fund, a grant funded by the the Wikimedia Foundation.

References

- Adafre, S. F.; and de Rijke, M. 2005. Discovering missing links in Wikipedia. In *Proc. LinkKDD '05*, 90–97. New York, NY, USA: ACM.
- Assmann, J.; and Czaplicka, J. 1995. Collective Memory and Cultural Identity. *New German Critique*, (65): 125–133.
- Bipat, T.; McDonald, D. W.; and Zachry, M. 2018. Do We All Talk Before We Type? Understanding Collaboration in Wikipedia Language Editions. In *Proc. OpenSym '18*, 1–11. New York, NY, USA: ACM.
- Bolikowski, 2009. Scale-free topology of the interlanguage links in Wikipedia. ArXiv:0904.0564 [physics].
- Buriol, L. S.; Castillo, C.; Donato, D.; Leonardi, S.; and Millozzi, S. 2006. Temporal Analysis of the Wikigraph. In *2006 IEEE/WIC/ACM Int'l Conf. on Web Intelligence*, 45–51.
- Consonni, C.; Laniado, D.; and Montresor, A. 2019. WikiLinkGraphs: A Complete, Longitudinal and Multi-

¹³English bot approval link: https://en.wikipedia.org/wiki/Wikipedia:Bots/Requests_for_approval

¹⁴Arabic bot approval link: <https://ar.wikipedia.org/wiki/مطلب>

- Language Dataset of the Wikipedia Link Networks. In *Proc. ICWSM '19*. AAAI.
- Dandala, B.; Mihalcea, R.; and Bunesco, R. 2012. Towards building a multilingual semantic network: identifying interlingual links in Wikipedia. In *Proc. SemEval '12*, 30–37. USA: ACL.
- Ferron, M.; and Massa, P. 2011a. Collective memory building in Wikipedia: the case of North African uprisings. In *Proc. WikiSym '11*, 114–123. New York, NY, USA: ACM.
- Ferron, M.; and Massa, P. 2011b. WikiRevolutions: Wikipedia as a Lens for Studying the Real-time Formation of Collective Memories of Revolutions. *Int'l J. Comm.*, (5).
- Ford, H. 2022. *Writing the Revolution: Wikipedia and the Survival of Facts in the Digital Age*. MIT Press.
- Gildersleve, P.; Lambiotte, R.; and Yasseri, T. 2023. Between News and History: Identifying Networked Topics of Collective Attention on Wikipedia. *Journal of Computational Social Science*, 6: 845–875.
- Halbwachs, M. 1992. *On Collective Memory*. Chicago: University of Chicago Press, 1st edition edition.
- Hale, S. A. 2014. Multilinguals and Wikipedia editing. In *Proc. WebSci '14*, 99–108. New York, NY, USA: ACM.
- Hara, N.; Shachaf, P.; and Hew, K. F. 2010. Cross-cultural analysis of the Wikipedia community. *JASIST*, 61(10): 2097–2108.
- He, S.; Lin, A. Y.; Adar, E.; and Hecht, B. 2018. The_Tower_of_Babel.jpg: Diversity of Visual Encyclopedic Knowledge Across Wikipedia Language Editions. In *Proc. ICWSM '18*, volume 12. AAAI.
- Hecht, B.; and Gergle, D. 2010. The tower of Babel meets web 2.0: user-generated content and its applications in a multilingual context. In *Proc. CHI '10*, 291–300. New York, NY, USA: ACM.
- Hickman, M. G.; Pasad, V.; Sanghavi, H. K.; Thebault-Spieker, J.; and Lee, S. W. 2021. Understanding Wikipedia Practices Through Hindi, Urdu, and English Takes on an Evolving Regional Conflict. In *Proc. CSCW '24*, volume 5, 34:1–34:31. ACM.
- Hoskins, A. 2017. *Digital Memory Studies: Media Pasts in Transition*. Routledge.
- Igarashi, N.; Okada, Y.; Sayama, H.; and Sano, Y. 2022. A two-phase model of collective memory decay with a dynamical switching point. *Sci Rep*, 12(1): 21484.
- Kansteiner, W. 2002. Finding Meaning in Memory: A Methodological Critique of Collective Memory Studies. *History and Theory*, 41(2): 179–197.
- Keegan, B. C. 2019. The Dynamics of Peer-Produced Political Information During the 2016 U.S. Presidential Campaign. In *Proc. CSCW '19*, volume 3, 33:1–33:20. ACM.
- Luyt, B. 2016. Wikipedia, collective memory, and the Vietnam war. *JASIST*, 67(8): 1956–1961.
- Massa, P.; and Scrinzi, F. 2012. Manypedia: comparing language points of view of Wikipedia communities. In *Proc. WikiSym '12*, 1–9. New York, NY, USA: ACM.
- Nemoto, K.; and Gloor, P. A. 2011. Analyzing Cultural Differences in Collaborative Innovation Networks by Analyzing Editing Behavior in Different-Language Wikipedias. *Procedia - Social and Behavioral Sciences*, 26: 180–190.
- Nora, P. 1989. Between Memory and History: Les Lieux de Mémoire. *Representations*, (26): 7–24.
- Pentzold, C. 2009. Fixing the floating gap: The online encyclopaedia Wikipedia as a global memory place. *Memory Studies*, 2(2): 255–272.
- Porter, E.; Krafft, P. M.; and Keegan, B. 2020. Visual Narratives and Collective Memory across Peer-Produced Accounts of Contested Sociopolitical Events. *Trans. Soc. Comput.*, 3(1): 1–20.
- Roy, D.; Bhatia, S.; and Jain, P. 2020. A Topic-Aligned Multilingual Corpus of Wikipedia Articles for Studying Information Asymmetry in Low Resource Languages. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, 2373–2380. Marseille, France: European Language Resources Association.
- Roy, D.; Bhatia, S.; and Jain, P. 2022. Information asymmetry in Wikipedia across different languages: A statistical analysis. *JASIST*, 73(3): 347–361.
- Rupprechter, T.; Burghardt, K.; and Helic, D. 2023. Poor attention: The wealth and regional gaps in event attention and coverage on Wikipedia. *PLoS One*, 18(11): e0289325.
- Twyman, M.; Keegan, B. C.; and Shaw, A. 2017. Black Lives Matter in Wikipedia: Collective Memory and Collaboration around Online Social Movements. In *Proc. CSCW '17*, 1400–1412. New York, NY, USA: ACM.
- Yasseri, T.; Gildersleve, P.; and David, L. 2022. Chapter 9 - Collective memory in the digital age. In O'Mara, S. M., ed., *Progress in Brain Research*, volume 274 of *Collective Memory*, 203–226. Elsevier.
- Zubrzycki, G.; and Woźny, A. 2020. The Comparative Politics of Collective Memory. *Annual Review of Sociology*, 46(1): 175–194.

Paper Checklist

1. For most authors...
 - (a) Would answering this research question advance science without violating social contracts, such as violating privacy norms, perpetuating unfair profiling, exacerbating the socio-economic divide, or implying disrespect to societies or cultures? **Yes**
 - (b) Do your main claims in the abstract and introduction accurately reflect the paper's contributions and scope? **Yes**
 - (c) Do you clarify how the proposed methodological approach is appropriate for the claims made? **Yes**
 - (d) Do you clarify what are possible artifacts in the data used, given population-specific distributions? **Yes**
 - (e) Did you describe the limitations of your work? **Yes**
 - (f) Did you discuss any potential negative societal impacts of your work? **Yes**
 - (g) Did you discuss any potential misuse of your work? **Yes**
 - (h) Did you describe steps taken to prevent or mitigate potential negative outcomes of the research, such as data and model documentation, data anonymization, responsible release, access control, and the reproducibility of findings? **Yes**
 - (i) Have you read the ethics review guidelines and ensured that your paper conforms to them? **Yes**
2. Additionally, if your study involves hypotheses testing...
 - (a) Did you clearly state the assumptions underlying all theoretical results? **N/A**
 - (b) Have you provided justifications for all theoretical results? **N/A**
 - (c) Did you discuss competing hypotheses or theories that might challenge or complement your theoretical results? **N/A**
 - (d) Have you considered alternative mechanisms or explanations that might account for the same outcomes observed in your study? **N/A**
 - (e) Did you address potential biases or limitations in your theoretical framework? **N/A**
 - (f) Have you related your theoretical results to the existing literature in social science? **N/A**
 - (g) Did you discuss the implications of your theoretical results for policy, practice, or further research in the social science domain? **N/A**
3. Additionally, if you are including theoretical proofs...
 - (a) Did you state the full set of assumptions of all theoretical results? **N/A**
 - (b) Did you include complete proofs of all theoretical results? **N/A**
4. Additionally, if you ran machine learning experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **N/A**
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **N/A**
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **N/A**
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **N/A**
 - (e) Do you justify how the proposed evaluation is sufficient and appropriate to the claims made? **N/A**
 - (f) Do you discuss what is "the cost" of misclassification and fault (in)tolerance? **N/A**
5. Additionally, if you are using existing assets (e.g., code, data, models) or curating/releasing new assets, **without compromising anonymity**...
 - (a) If your work uses existing assets, did you cite the creators? **N/A**
 - (b) Did you mention the license of the assets? **N/A**
 - (c) Did you include any new assets in the supplemental material or as a URL? **N/A**
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? **N/A**
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? **N/A**
 - (f) If you are curating or releasing new datasets, did you discuss how you intend to make your datasets FAIR? **N/A**
 - (g) If you are curating or releasing new datasets, did you create a Datasheet for the Dataset? **N/A**
6. Additionally, if you used crowdsourcing or conducted research with human subjects, **without compromising anonymity**...
 - (a) Did you include the full text of instructions given to participants and screenshots? **N/A**
 - (b) Did you describe any potential participant risks, with mentions of Institutional Review Board (IRB) approvals? **N/A**
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? **N/A**
 - (d) Did you discuss how data is stored, shared, and deidentified? **N/A**