

Analysing Timelines of National Histories across Wikipedia Editions: A Comparative Computational Approach

Anna Samoilenko,^{1,2} Florian Lemmerich,^{1,2} Katrin Weller,¹
Maria Zens,¹ Markus Strohmaier^{1,2}

¹ GESIS – Leibniz-Institute for the Social Sciences

² University of Koblenz-Landau

Abstract

Portrayals of history are never complete, and each description inherently exhibits a specific viewpoint and emphasis. In this paper, we aim to automatically identify such differences by computing timelines and detecting temporal focal points of written history across languages on Wikipedia. In particular, we study articles related to the history of all UN member states and compare them in 30 language editions. We develop a computational approach that allows to identify focal points quantitatively, and find that Wikipedia narratives about national histories (i) are skewed towards more recent events (*recency bias*) and (ii) are distributed unevenly across the continents with significant focus on the history of European countries (*Eurocentric bias*). We also establish that national historical timelines vary across language editions, although average interlingual consensus is rather high. We hope that this paper provides a starting point for a broader computational analysis of written history on Wikipedia and elsewhere.

1 Introduction

Writing about history – historiography – is important in all social groups. Establishing some consensus on relevant dates provides a feeling of roots, and is at the core of building identities – for individuals, groups, or nations. Each description inherently presents a unique viewpoint on past events, and it might be partial and disputable. Today, the online encyclopedia Wikipedia has a vast readership across continents and languages. It offers quick, effortless access to a spectrum of reference information, including historical accounts. These representations might contain errors and false information (Potthast, Stein, and Gerling 2008), be biased towards specific viewpoints, or differ otherwise from the books written by professional historians. Fortunately, Wikipedia’s open and digital nature allows for thorough quantitative analysis of historical narratives, even across a large number of languages – something which is not a typical case for other historiographical sources, such as printed encyclopedias or history textbooks.

In this paper, we investigate descriptions of national histories in different Wikipedia language editions, taking a comparative computational approach. In that direction, we pursue two goals: (1) presenting a data-driven approach that

enables analysis of historiography through a computational lens, and (2) answering specific research questions on the depiction of history in Wikipedia.

Approach. We present a computational approach to the analysis of textual historiographical data which is suited for large-scale comparative studies. We apply it to Wikipedia articles on all UN member states in 30 language editions. We concentrate on *year dates* as accessible representations of more complex historical structures. To be able to compare the descriptions across languages, we retrieve from article texts all date mentions (in the form of 4-digit numbers between 1000-2016), and use them as a language-independent unit of comparison (Rüsen 1996). We propose a simple randomisation technique to extract *significant focal points of national histories* – time periods of significantly high mentions, compared to a random expectation model. We combine visual interpolation and expertise of history experts in order to evaluate how the results of our approach compare with the existing historical knowledge. We use hierarchical clustering to group countries whose histories are represented similarly on Wikipedia. Finally, we compute inter-language agreement on history of each country using the Jensen-Shannon divergence measure.

Empirical questions. We compare, what readers of different languages can learn about national histories from their ‘home’ Wikipedia language editions. In particular, we focus on three research questions: **RQ1:** What are the most documented periods of history of the last 1,000 years in Wikipedia? **RQ2:** What are the temporal focal points in descriptions of national histories in Wikipedia? **RQ3:** Are country timelines consistent across language editions?

Empirical findings. We find the presence of recency bias across all language editions and countries – the tendency to document recent events more frequently than those that happened in a more distant past. We also find that the distribution of historical focal points in the analysed articles is inhomogeneous across continents. We discover a multitude of focal points distributed through entire timelines of European countries, while we see much fewer highlights in pre-Columbian Americas and Oceania. Groups of countries with similarly distributed focal points map well to geopolitical blocs. Finally, we find differences in the way national histories are described in the examined language editions, although on average the cross-lingual consensus is rather high.

Contributions. We contribute to the pool of computational methods that help to quantify historiographical processes. We combine multiple computational methods into an approach that can be used for quantitative analysis of large textual historical and historiographical datasets, such as demographic and economic records, census data, digitised books, etc. Our approach scales well, and is suited for large comparative studies of multidimensional (e.g. multiple languages and countries) data.

By including 193 countries and 30 languages, we step beyond the current state of comparative historiography and allow for a large-scale transnational perspective on similarities, conjunctions, or alternatives in historiography. Although we start from the (limited) concept of the (pre-) histories of nation-states we, (i) methodologically, enable cross-lingual and -national clustering and comparison and, (ii) empirically, show that historiographical focal points transcend national borders, and contribute to existing literature on collective memory and public history as created and perceived through Wikipedia.

2 Related literature

Our approach carries characteristics of ‘the digital turn’ that the study of history has envisaged in the recent years: it uses a large (digital) data set, borrows from statistical methods, and conceptually, turns to transnational and global comparative perspectives. Theoretically, our approach lies in the domain of cultural history (analysing the multitude of historical interpretations, e.g. gender-based or post-colonial histories), with a specific focus on the formation and effects of collective/public memories (Conrad 2007), and the analysis of nations as imagined communities (Anderson 2016).

Wikipedia as a data source: Many non-academics start seeing history as a venue for active participation, rather than a domain of professional historians (Rosenzweig and Thelen 1998). The encyclopedia Wikipedia is open for everyone to contribute on any topic. Thanks to this feature, it has become one of the primary venues where ‘free-lancer’ amateur historians, potentially, alongside with professionals, can participate in history-making and shaping historiographic discourse (Conrad 2007). A number of professional historians recognise Wikipedia as a place where enthusiasts collaboratively re-think the past (Pfister 2011), construct public memories (Pentzold 2009), and write history in an open source manner (Rosenzweig 2006). Although such popular understandings of the past might differ from those of professional historians (Conrad 2007), Wikipedia is a popular source of information when it comes to history (Spoerri 2007) and thus has become an object of research itself.

To the best of our knowledge, only a few researchers have investigated historical narratives of Wikipedians: Luyt (2011) compared the articles on the history of two countries, concluding that the history of Singapore recounts the dominant political narrative, while the article on the history of the Philippines contains both traditional and alternative views. Jensen (2012) looked into the discussion pages about the article on the war of 1812 and found that the main debate among the editors was on who won the war. Both studies use a traditional descriptive methodology. Finally, Gieck et

al. (2016) used a data science approach and compared war-related articles across five language editions, using methods from sentiment-, network-, and language complexity analysis. The authors found that World Wars I and II are the most important historical events in these editions.

Quantifying history: Quantitative approaches were integrated into history studies in the last century. Opposite to traditional qualitative interpretations, they relied on statistical methods and a new conceptualisation, in which historical reality was condensed to quantifiable (often socioeconomic) historical facts, whose evolution was traced through longitudinal studies (Furet 1971). Computational approaches that allow formulating and testing retrospective hypotheses, running historical experiments, and discovering large-scale patterns of the past by processing big data-sets, appear to be the obvious next step that could turn history into an analytical, deductive, predictive science (Turchin 2011; Kiser and Hechter 1991). Mathematical simulations have helped to test historical hypotheses about the evolution of commodity flows across ancient Asia (Malkov 2014), the influence of agriculture on birth rates in the Old World (Bennett 2015), and the rise and fall of large-scale societies (Turchin et al. 2013). Network approaches have also found popularity among historiographers, due to their possibility to visualise and quantify relational ties that are abundant in written (digitised) historical texts. Networks have been used to map the travels and settlement of the Vikings (Sindbæk 2007), to study centralisation of political parties in Renaissance Florence (Padgett and Ansell 1993), to examine interconnections between the elite individuals of Medieval Scotland (Jackson 2016), and to track the intellectual mobility of notable individuals on a large scale (Schich et al. 2014).

The step from statistic to historic interpretation still remains a difficult one, which is one of the reasons quantitative approaches have been slow in gaining support among the traditional historians. Nevertheless, computational studies could contribute evidence to support the existing historical theories, or suggest otherwise unavailable new hypotheses. The recent rise of interest to digital humanities and successes in digitising large collections of historical documents (Michel et al. 2011; Yu et al. 2016) allow broad possibilities for historians to select the data relevant to their questions. Still, the pool of available methods remains rather sparse. As it becomes easier for historians to extract data from digitised records, as well as from digitally born sources such as Wikipedia, new methods are in need that will help quantify and map historical processes.

3 Data collection & validation

In this section we describe the steps of data collection and validation. Both data and code are available online (Samoilenko 2017). We focus on the history of 193 countries¹ which are the current UN member states².

¹Throughout the paper we use the terms nation, country, and state as synonyms, being aware of the differences.

²List of the UN member states, <http://www.un.org/en/member-states/index.html> (accessed Nov. 13, 2016)

Table 1: Expected error rates of dates extraction: language editions and centuries

Language	Exp. error.	Language	Exp. error.	Language	Exp. error.	Century	Exp. error.
English	0.0067	Ukrainian	0.0351	Basque	0.0042	11th century	0.2428
German	0.0186	Catalan	0.0096	Bulgarian	0.0262	12th century	0.0442
Swedish	0.0109	Norwegian Bokmal	0.0538	Danish	0.0086	13th century	0.0982
Dutch	0.0198	Serbo-Croatian	0.0068	Slovak	0.0040	14th century	0.0214
French	0.0018	Finnish	0.0208	Lithuanian	0.0266	15th century	0.0363
Russian	0.0395	Hungarian	0.0347	Croatian	0.0178	16th century	0.0415
Italian	0.0081	Romanian	0.0378	Slovenian	0.0025	17th century	0.0261
Spanish	0.0069	Czech	0.0076	Estonian	0.0131	18th century	0.0089
Polish	0.0223	Serbian	0.0119	Galician	0.0154	19th century	0.0094
Portuguese	0.0166	Turkish	0.0256	Norwegian Nynorsk	0.0246	20th century	0.0000
						21st century	0.0100

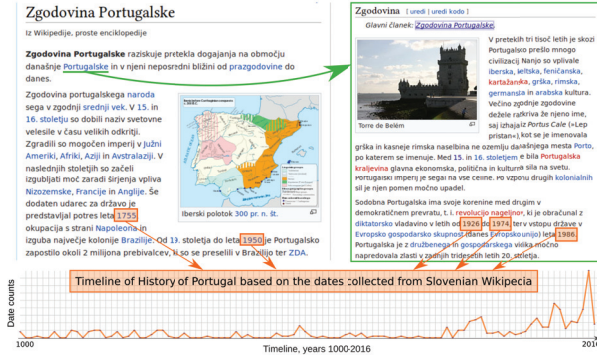


Figure 1: Data collection. We show parts of the article on Portuguese history and one of its outlinks, as they appear in Slovenian Wikipedia in 2016. We collect all 4-digit numbers from the main text of the article and all its outlinks, and analyse the resulting distribution (bottom part of the figure).

Data collection: For each of the 193 countries we locate an article in the English edition of Wikipedia, titled 'History of X', where X is the country name. Using Wikipedia's inter-language links, we retrieve other language versions of the article from sister editions. We limit the analysis to 30 largest Wikipedia editions (more than 125,000 articles³) providing these languages are native to Europe. By applying this setup we avoid issues connected with extraction and alignment of dates from the languages using different calendars and alphabet systems. The limitations related to multilingual data retrieval and the choice of linguistic scope are discussed in detail in Section 6.

We retrieve the main text of each article (as HTML, excluding text related to, e.g. references) from the English Wikipedia and – if available – from all 30 selected sister editions, as well as the text of all Wikipedia articles to which these pages link. We focus on the out-links because they are embedded in the main articles' texts and thus immediately available for a reader to inspect, unlike, for example, the in-links, which could not be found by reading the article page. We find between 14,927 (Italy) and 394 (the Federated States of Micronesia) articles related to the history of each

nation. In order to assess the coverage of historical periods, we choose a language-independent measure – the mentions of year numbers in the article text. Since we are interested in historical events of the last millennium, we retrieve all 4-digit numbers in the range between 1000 and 2016 from the main text of all articles in our collection. Fig. 1 illustrates the process with an example of an article on the history of Portugal in Slovenian Wikipedia. In cases when paragraphs consist mostly of hyperlinks (more than 50% of words are hyperlinks), we record no dates from them, since there is little narrative in such paragraphs.

We ran the data collection in July 2016, using the access provided by Wikimedia Tool Labs⁴ as well as the Wikipedia API⁵. We retrieved approximately 17M dates from 773,121 articles in 30 language editions.

Data validation: In order to ensure internal reliability of our extraction method, we check whether the extracted numbers are years rather than numerals indicating, for example, height. We create a random sample of 3,300 extracted 4-digit numbers evenly split across 30 languages and 11 centuries, and ask 3 independent human coders to evaluate each case, i.e. to say whether a number is a date or not (false positive). For each language there are 110 evaluation tasks, which consist of: the potential date (4 digits), the text surrounding the potential date in the original language (40 characters before and after the number), and the same text translated into English via Google Translate (except for the English edition case). If the coder is unsure about a number, we treat it as a false-positive. Each case is settled by the majority vote. The inter-rater agreement is substantial (Fleiss' kappa = .77). We compute the expected error rates for centuries,

$$\langle E_c \rangle = \frac{1}{D_c} \sum_l \left(\frac{n_{c,l}}{10} d_{c,l} \right), \quad (1)$$

and language editions,


$$\langle E_l \rangle = \frac{1}{D_l} \sum_c \left(\frac{n_{c,l}}{10} d_{c,l} \right), \quad (2)$$

where D_l and D_c are the total count of collected potential dates per language and century, and $n_{c,l}$ is false positives


⁴Wikimedia Tool Labs, <https://wikitech.wikimedia.org/> (accessed Nov. 13, 2016)

⁵Wikipedia API for Python, <https://pypi.python.org/pypi/wikipedia/> (accessed Nov. 13, 2016)

³Wikipedia: List of Wikipedias, http://en.wikipedia.org/wiki/List_of_Wikipedias (accessed Nov. 13, 2016)



(a) Extracted timelines: 30 language editions.



(b) Extracted timelines: selected countries

Figure 2: Distribution of collected dates. (a) in 30 language editions and (b) in top 10 countries according to the number of collected dates. Across editions of all sizes, and across countries we observe the same strong bias towards dates within the last 100 years, while dates from before 1500 are rarely mentioned.

count in our random sample of $d_{c,l}$ numbers collected per language edition l and century c .

We report the expected error rates for both centuries and language editions in Table 1. All language editions and most of the centuries show a very low expected error rate (below .04). We estimate the highest error rate in the 11th century (.24), since a large number of extracted digits turned out to be numerals relating to heights, population counts, etc. This error is present mostly in this century, presumably due to the numeral 1000 being often used for other purposes than mentioning a year. Other false-positives include dates from Before Christ. In the more recent centuries our extraction method is very exact.

4 Approach and results

In this section, we describe our approach and present the findings regarding inter-lingual portrayal of national histories in Wikipedia. The approach consists of three parts: 4.1 identifying most covered historical periods across countries and languages, 4.2 extracting the focal points of national histories across all language editions, and 4.3 quantifying the amount of inter-language agreement on representation of national histories.

4.1 Most covered historical periods

In order to gain a better understanding of our dataset, we first look into timelines of extracted dates. We present the distribution of collected dates across all 30 language editions (Fig. 2a) and ten selected countries (Fig. 2b) – those with the most available dates.

Results. Across language editions, the data show a bias towards recent dates, having a large proportion of dates (between 60 and 80 percent) in the more recent decades (since 1800), and very low date counts before 1500. This is partly due to the chosen subject (nation states), but also points to a more general recency bias. Thus, Wikipedia readers can find a more detailed documentation of historic events of the past 200 years, compared to earlier centuries. Apart from intense coverage of the most recent events (2000s), we also observe peaks of date mentions that correspond to some of the most

violent recent conflicts: Napoleonic war (1800-10s) and the First (1910s) and the Second (1940s) World Wars.

4.2 Historiographic focal points of countries

We build a Null Model that describes expected frequencies of dates in each decade under the assumption that all countries are presented equally. By comparing the observed frequencies with the expected baseline, we are able to detect significant focal points. For this part of the analysis we aggregate all country-related dates across language editions.

Method: A Null Model of focal points. Our approach is essentially an urn model, with replacement and without duplication, based on the underlying Dirichlet-multinomial distribution. First, we create a pool M of the dates found in all language editions for all countries. Then, for each country i , we randomly draw from the pool N_i dates, where N_i is the number of dates related to the history of country i collected from all language editions. We then count how many of these dates fall in each decade. We repeat the process 1,000 times. For each decade we thus build a distribution of the number of dates it can contain, within the hypothesis of events randomly distributed in time. Furthermore, we can compare the mean of the expected distribution $E[w_i^d]$ with the empirical date count for the country in the same decade, w_i^d , and convert the difference into a z -score. This procedure allows us to identify for each country in which decades the number of observed dates w_i^d differs significantly from the expected number of dates in this decade. Mathematically, the z -score of country i in decade d is given by:

$$z_i^d = \frac{w_i^d - E[w_i^d]}{\sigma_i^d}, \quad (3)$$

where $E[w_i^d]$ is the mean of the simulated date counts in decade d across 1,000 random draws, and σ_i^d is the standard deviation of the simulated date counts.

Fig. 3 illustrates the method with a toy example. We consider four hypothetical countries A-D with different artificial timelines (Fig. 3a). The dates are binned into decades, each matrix cell indicates the number of dates in the corresponding decade. Fig. 3b shows histograms of date counts

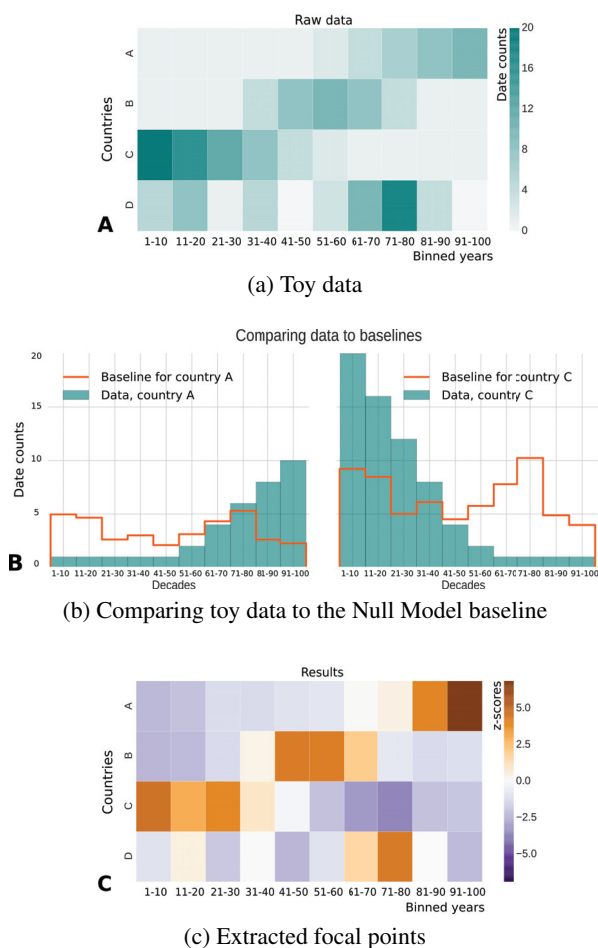


Figure 3: Illustration of the method. Fig.3a shows the initial distribution of toy data for hypothetical countries A-D. The dates are binned into decades, each cell colored according to the number of dates. Fig. 3b illustrates the method on countries A and C. We plot a histogram of date counts for each country (green bins), and compare them with expected baselines (orange lines). The baselines are the average over four initial distributions, adjusted to match each country’s total date count. Thus, the baselines are unique for every country and decade. Finally, in Fig. 3c we convert the differences between the data and the baseline into z -scores.

per country. Orange lines correspond to expected distributions given by the Null Model, which are simply the average over four initial distributions, adjusted to match the total country date count. These baselines vary across decades and countries so that to account for differences in countries’ total datecount. We then convert the differences between the observed data and the baselines into z -scores (Fig. 3c). For each country we can now extract the decades in which the number of collected dates differs significantly from the expected baseline. This method is especially useful when differences in counts are not directly comparable, as we have in our case: all row sums in Fig. 3a differ. To illustrate, in Fig. 3a the cell with the largest datecount is country C in the first

decade, but after comparing with the expected baseline (Fig. 3b), country A in the last decade stands out most, although its underlying count is smaller.

Results. The results of this procedure for selected geopolitical blocs are reported in Fig.4. z -scores below -6 and above 6 correspond to Bonferroni-corrected p -values < 0.01 (the expected distributions of decade counts are approximately normal), which means the results those cells are statistically significant. There are noticeable differences in distributions of focal points (in dark orange) across countries. For Western European countries, we observe high coverage of the Medieval and Early Modern periods (until ~ 1800). Specific periods of interest for individual countries include, for example, the French Revolution in France (1780-90s) and the Third Reich in Germany (1930-40s). By contrast, in East Asia the focal points are more heterogeneous. For Mongolia, the timeline focuses on the Mongolian Empire in the 13th century. Articles on Japanese and Chinese histories exhibit a strong focus on specific small time frames: the rise of the Tokugawa shogunate (1180-90s), the Kenmu Restoration (1330s) and the beginning of the Edo period (around 1600) in Japan; and the rise of the Jin (1120s), Yuan (1270s), Ming (1360s) and Qing (1640s) dynasties in China. Only with stronger European involvement in the region (starting in the mid-19th century) there is a more steady coverage. For Central America, the timelines focus on the Age of Discovery (late 15th - early 16th centuries), and the Spanish-American Wars of Independence (first half of the 19th century). In North America, the eras of the American Revolutionary War (end of 18th century) and the American Civil War (1860s) are most noticeable. For different regions of Africa, historical timelines strongly focus on the periods of its occupation and colonisation (Scramble for Africa in late 19th century), and recent history following its decolonization in the 1960s. In contrast to Southern Africa, North African national timelines focus on the Medieval history (Caliphate era), which is also the time of close interaction with Europe. The coverage seems to seize around 1300, just before the outbreak of the Black Death epidemic. For Australia and New Zealand the peaks in 1760-80s correspond to the expeditions of James Cook discovering Oceania and South Pacific. Over the next centuries, as contacts between Europeans and the local population grew, the coverage remains stable.

Overall, the number of discovered ‘focal points’ differs across regions. Within 30 examined Wikipedia editions, there is a disproportionate focus on histories of European countries, and the coverage of non-European states seems more intense in the periods when those states had closer interactions with Europe.

Clustering. We use the results reported in Fig.4 to group countries based on their historical timelines’ similarity. We represent each country as a vector of z -scores, and group together the countries whose z -score values across decades are similar both in direction and intensity. We compute pairwise cosine similarity between all countries, and apply hierarchical clustering with complete linkage (Müllner 2011) on top of the obtained values. The resulting dendrogram shows that the clusters correspond rather well to geopolitical regions. To illustrate this point, we cut the dendrogram at an




Figure 4: Temporal focal points of selected countries⁶ z-scores below -6 and above 6 correspond to Bonferroni-corrected p -values < 0.01 , which means the results in all coloured cells are statistically significant. Higher z-scores (orange) correspond to positive differences between the observed and the expected date count per decade. Cells with fewer than 30 dates are masked out. Interpretation of historical events corresponding to some focal points (in green) is offered by history experts. The distributions of focal points suggests there are similarities across countries within geopolitical blocs.

(arbitrary) level $t=2$, and plot the resulting 18 country clusters on a world map (Fig. 5). It suggests that focal points of the countries from the same geopolitical regions are similar in the editions that we analyse. It is important to underline, however, that similar focal points do not necessarily indicate a reference to the same historical events.

Combining information about cluster membership in Fig. 5 with the significant focal points presented in Fig. 4, we get a transnational impression on patterns of similarity among national timelines. We see, for example, that most of Africa maps to one cluster. Despite individual differences between country histories, in the analysed descriptions, history of the entire continent is reduced to the periods of its (de-)colonisation. Similarly, focal points of most of Central and South American countries are limited to the Age of Discovery and their Wars of Independence. On the contrary, Europe is separated into several clusters, as here the differences among the individual national timelines are more distinct. This gives an impression of how the entire world groups into regions based on the extracted focal points of individual countries. Also, it illustrates that for some parts of the world (e.g. Africa and parts of Americas), the analysed timelines show a reduced view of history.

⁶Additional plots illustrating the distribution of focal points across all countries, and the complete dendrogram of hierarchical clustering could be viewed at <http://annsamoilenko.wixsite.com/homepage/historical-landscapes>

4.3 Quantifying inter-edition agreement

In this section we investigate if national historical timelines are consistent across languages. For that, we compute a measure of their divergence across Wikipedia editions.

Method: Jensen-Shannon divergence. Based on the extracted probability distributions of years, for each country we compute a matrix of pairwise inter-language dissimilarities, using the Jensen-Shannon (J-S) divergence (Lin 1991):

$$J(p \parallel q) = \frac{1}{2} \left[\sum_t p(t) \log_2 \left(\frac{2p(t)}{p(t) + q(t)} \right) + \sum_t q(t) \log_2 \left(\frac{2q(t)}{p(t) + q(t)} \right) \right], \quad (4)$$

where $p(t)$ and $q(t)$ refer to the probability of year t in the language editions p and q . The divergence $J(p \parallel q) \in [0, 1]$, with 0 indicating complete overlap between the compared distributions. Differences across language-specific timelines of each country are summarised by a square $m \times m$ matrix, where m is the number of extracted language editions (up to 30) covering the country's history. We summarise interlingual differences in two values: median and spread of the distribution of $J(p \parallel q)$ values per country, and present them in a scatterplot for ease of visual analysis.

Results. The results of this approach are presented in Fig. 6. Data points in the lower left quarter correspond to countries with the lowest medians and the narrowest distri-

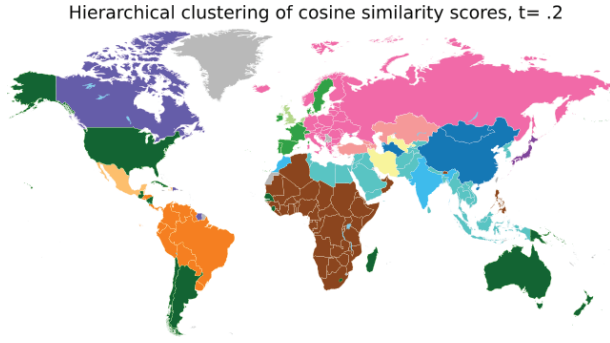


Figure 5: World map of country clusters⁶. This results from cutting the dendrogram of hierarchical clustering at a threshold $t = .2$. The countries within coloured clusters have similar temporal focal points, based on the articles in analysed editions, and map well to geopolitical regions.

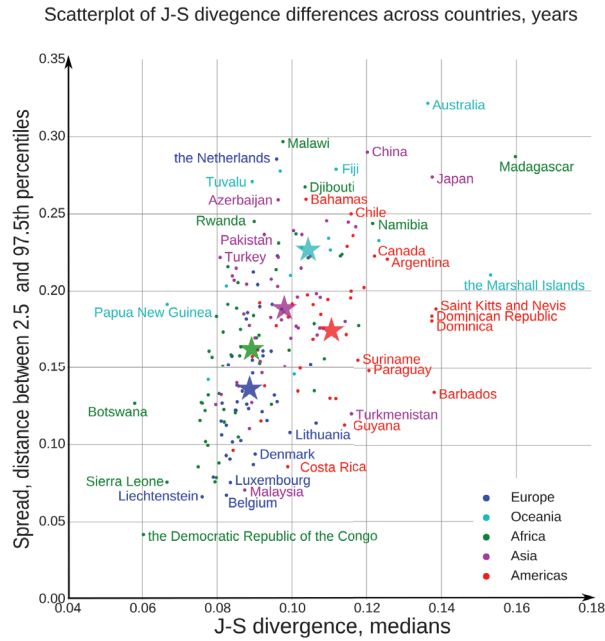


Figure 6: Inter-language consensus in Wikipedia articles on national histories, based on pairwise Jensen-Shannon divergence values. Countries in the lower left of the plot show the highest consensus across editions. Stars represent data centroids for countries of the same region. The plot shows a high inter-language consensus on average, though the descriptions are not identical across editions. European countries exhibit the highest amount of consensus.

butions of J-S scores (i.e. the smallest differences between the most similar and the most different language pair), and thus, with the highest inter-lingual consensus. Overall, J-S scores are centered around very low values (medians between .06 and .16), which indicates a high average agreement across language editions. Their spread covers a higher range (up to .35), implying the presence of large differences between some language pairs. Based on the location of data

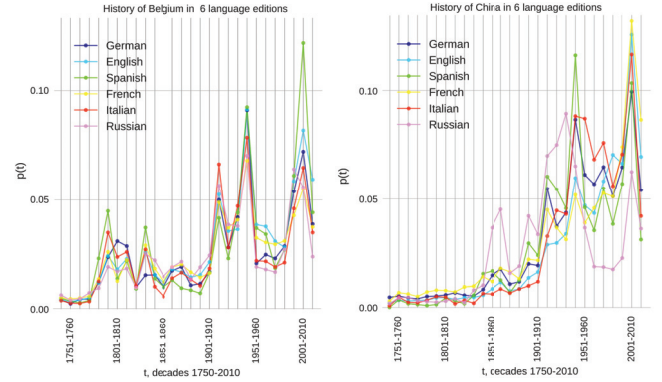


Figure 7: Inter-lingual consensus on histories of selected countries. To illustrate cases with very high (a) and very low (b) inter-lingual consensus, we present parts of probability distributions of dates zoomed into 1750-2010s, for 6 large editions. Chinese timeline in Russian Wikipedia differs noticeably from the timelines in all other editions, while for Belgium all timelines are almost identical.

centroids (stars), we observe higher interlingual consensus on the history of European and African countries, compared to Americas, Asia, and Oceania. The largest interlingual disagreement is found in the articles on the history of Australia, Malawi, Madagascar, China, Japan, and the Netherlands; some with the highest consensus are Liechtenstein, Belgium, the Democratic Republic of Congo, and Malaysia. In case of China (Fig. 7b), for example, high disagreement is partially driven by the differences between Russian and other language editions. This is especially evident during the period of Sino-Soviet split in 1960-80s, which is less densely covered in the Russian language Wikipedia. Timelines of history of Belgium, on the other hand, are almost identical across all 30 language editions (in Fig. 7a we present only 6 largest editions in order not to obstruct the view). Overall, we conclude that country-specific historical timelines differ across language editions, although on average such differences are rather small.

5 Discussion

In this section we reflect on characteristics of the proposed computational approach, and discuss our empirical findings.

Approach: We have selected year mentions as a language-independent unit of analysis that can act as a measure of emphasis on certain historical events (this formalisation has been used before by Michel et al. (2011)). In principle, it is possible to choose any other quantifiable unit, such as mentions of geographical locations, persons, or events, providing there is an extra step for tackling inter-lingual entity disambiguation. Although the building blocks of our approach are not new in computational fields, they are new to the community of quantitative historians, where hitherto unseen inflow of newly digitised historical records pose a

methodological challenge. Our approach is general enough to be applied to many large digitised datasets (such as demographic and economic records, census data, books, etc.), and it is suitable for comparative analysis of any number of countries across languages. By applying purely computational and data-driven methods, we are able to eliminate the bias that could be posed by the researcher’s cultural background (Ailon 2008), and perform transnational analysis on a scale previously unknown to comparative historiography.

Empirical results: We apply our approach to a case study and investigate, what readers of different languages can learn about national histories from 30 Wikipedia editions. Our empirical results indicate the presence of *recency bias* across language editions and countries: most retrieved dates belong to the recent decades, while those before 1500 are very sparse. Other studies report similar findings: a survey of students from Europe, the US, and Japan about the events that they perceived most important in the last 1,000 years showed that 60% of the mentioned events happened in the last 300 years (Rovira, Deschamps, and Pennebaker 2006), while in our case it is between 60 and 80% depending on the language edition. Recency bias is a well-known concept in the fields of social/collective memory and psychology of history, and it is sometimes referred to as genealogical (Candau 2005), autobiographical (Wertsch 2002), or, most commonly, communicative memory (Assmann 2011; Assmann and Czaplicka 1995). We find that recent memories are actively documented on Wikipedia; possibly, because we simply know more about the recent past. However, like (Pentzold 2009), we also find evidence that these narratives stretch beyond the limited domain of communicative memory (‘floating gap’ of 3-4 generations), and reflect long-term, stabilised cultural memory (Assmann 2011). While recency bias is common in oral accounts of history, it is novel to demonstrate it for the context of a written encyclopedia.

The analysis of historiographic focal points of countries indicates inhomogeneous coverage across the continents, but high similarity within geopolitical country blocs. We find a multitude of focal points distributed across the whole timelines of European countries, while we observe very sparse coverage and no focal points in pre-Columbian Americas and Oceania. Significant focal points in non-European states appear to relate to the periods which are culturally and historically important for Europe, such as the discovery of Latin America and the Polynesian islands by European travellers, the beginning of European trade with China, and the period of close interaction between Europe and Northern Africa up until the Black Death epidemic. We interpret this as evidence of *Eurocentric bias*, an issue well-documented in professional historiography (Geyer and Bright 1995), but also present in public perceptions of history, as cross-cultural surveys show (Rovira, Deschamps, and Pennebaker 2006; Liu et al. 2005). Thus we find that Wikipedia, despite offering a democratised way of writing about history, reiterates similar biases that are found in the ‘ivory tower’ of academic historiography. Given that the focus of our analysis is on languages spoken in Europe, some dominance of Eurocentric perspectives is expected. Still, some languages that we study (e.g., English and Spanish) are widely spoken

in other regions, such as Latin America and Africa. Considering their international reach and the collaborative, global nature of Wikipedia, it is surprising to empirically confirm this imbalance towards European countries.

We also find *high consensus across the examined editions* in describing individual country histories. Across language editions, extracted dates also peak in the same decades, which correspond to periods of highly violent conflicts. Although this finding is not immediately intuitive, previous studies have reported high consensus in how different cultures view important historical events (Rovira, Deschamps, and Pennebaker 2006). The authors explained this by the possible existence of cross-cultural collective memory, dominant hegemonic beliefs about the world history, and the narrowing cultural and interest differences between the communities. The latter has also been studied in Wikipedia context, finding that both linguistic (Samoilenko et al. 2016) and geographic communities (Karimi et al. 2015) of Wikipedia editors are interested in similar article topics. We add to this research by demonstrating that in the case of history, the content (recovered timelines) of articles is also very similar across languages. These similarities in the public perceptions of history might be a result of a converging approach to history education, common exposure to media and entertainment (such as popular history TV shows), or lack of exposure to alternative historical material (Conrad 2007). Additionally, cross-lingual Wikipedia editors and bots might be responsible for inserting similar material in different language versions of the article (Hale 2014), which might be a factor in the similarities we find.

6 Limitations

Our methodological choices have implications on the conclusions we can draw from the analysis. Below we present the limitations of this study grouped by type.

Unit of analysis: One of the main limitations of any computational approach is the fact that it is reductional: in this study, we reduce the complexity of historiography to year mentions. The advantage of this approach is in focusing on an objective, quantifiable unit that can be compared across language editions without the biases introduced by translation. We leave other, language-specific methods of analysing historical narratives and sentiments outside of the scope of this study. Our measure, a year, is a very fine-grained unit when examining a millennium of human history. Still, by leaving mentions of decades or centuries unaccounted for, we lose some precision, especially for certain countries and earlier time periods where the exact dates of events are not well-known or documented.

Focal points: We define focal points as time periods of significantly high mentions, compared to a random expectation model. Other formulations of Null Models are possible, which could describe a random process otherwise, and potentially result in non-identical outcomes. However, our model makes the least assumptions about the countries’ histories and thus makes a good baseline. Interpretations of historical events related to some extracted focal points depict a viewpoint of selected history experts and are subjective.

Linguistic scope: We focus on 30 largest editions of Wikipedia which are also the languages native to the geographic region of Europe. For these editions, year is an acceptably robust unit of analysis, since these languages generally share date and time notation standards. We exclude other large editions such as Chinese, Arabic, and Farsi since their distinctive calendar- and numeral systems require developing language-specific methods of dates extraction, and this task goes beyond the scope of this study. The conclusions of our study should not be generalised to the whole Wikipedia, and are only valid for the studied editions.

Multilingual data retrieval: Our method of retrieving sister-articles from non-English language editions relies on Wikipedia's inter-language links (ILLs). Although the quality of ILLs is a debatable issue, studies have shown that the proportion of bidirectional ILLs between English and the largest European languages is around 98% (Rinser, Lange, and Naumann 2013). We do not exclude the possibility that some of the multilingual articles might have been missed, however it is reasonable to assume that their absolute share will not have dramatically affected the results of the study.

Article disambiguation: Our analysis focuses on the articles with the specific title wording, 'History of X'. To solve title disambiguation issues in the English edition, we manually map all countries to corresponding Wikipedia articles on their modern history. In cases when a territory has changed names several times (having been a part of several countries, e.g., post-Soviet bloc), there might be multiple Wikipedia articles related to its history. We partially tackle this issue by including out-linked articles in our dataset, which in itself may introduce additional noise, and may potentially impact the findings.

Inhomogeneous data: We compare language editions at different ages, states of saturation, and sizes of underlying potential editor populations. This unavoidably leads to an overrepresentation of larger editions, for example, the pool of dates in RQ2 (Section 4.2) is heavily influenced by German and English editions.

Data validity: Data validation has shown high accuracy of our date extraction method. This is possible because we limited our analysis to the articles evidently related to history. The precision of the method might suffer when analysing texts of broader scope or focusing on the dates from Before Christ era. Already in our sample, we find small numbers of false-positives, e.g. 4-digit numerals expressing heights, lengths, or population counts. Although suitable for the current setup, our dates extraction method might need improvement if applied to a different dataset.

7 Conclusions

Our study explores Wikipedia as a novel data source for the community of historiographers. We focus on multilingual Wikipedia articles about country histories, and study popular perceptions of the past which are collaboratively created by non-professional history enthusiasts.

For this study, we have developed a computational approach which is scalable, language-independent, and flexible in terms of choosing its unit of analysis. We have used

the approach to extract and visualise the periods of significant historical importance of 193 countries over the last 1,000 years, based on Wikipedia articles about their history in 30 language editions. We find that public narratives about history on Wikipedia are skewed towards more recent events, and are distributed unevenly across the continents, with significant focus on the history of European countries. We also find that cross-lingual consensus on national historical timelines is on average rather high, although disagreements are also present.

The observed 'peaks' and 'lows' of interest to certain time periods, as well as cross-lingual differences in national timelines, might have different explanations. If these dissimilarities are *intentional*, they might be a reflection of cultural differences, and in this case, our results could be interesting for historians and cultural scientists who might wish to explore the topic in greater detail and with other methods. If these differences are *accidental* or could be reduced to 'missing data', our findings could be actionable for the Wikimedia community and enthusiastic editors who wish to improve the quality of the articles in various language editions. In any case, our results show that Wikipedia's historical reference articles are not free from gaps and biases. We hope that History teachers and students, as well as lay readers who use Wikipedia to enrich their knowledge about world history, would benefit from this awareness.

8 Acknowledgements

The authors would like to express their sincere gratitude to the members of GESIS CSS Data Science team, and in particular, to Mohsen Jadidi, Fariba Karimi, Mathieu G  nois, and Sebastian Stier for insightful discussions of the project during its earlier stages.

References

- Ailon, G. 2008. Mirror, mirror on the wall: Culture's consequences in a value test of its own design. *Academy of Management Review* 33(4):885–904.
- Anderson, B. 2016. *Imagined communities: Reflections on the origin and spread of nationalism*. Verso London.
- Assmann, J., and Czaplicka, J. 1995. Collective memory and cultural identity. *New German Critique* (65):125–133.
- Assmann, J. 2011. Communicative and cultural memory. In *Cultural Memories*. Springer. 15–27.
- Bennett, J. 2015. Modeling the large-scale demographic changes of the Old World. *Cliodynamics: The Journal of Quantitative History and Cultural Evolution* 6(1).
- Candau, J. 2005. *Anthropologie de la m  moire*. Armand Colin.
- Conrad, M. 2007. 2007 Presidential Address of the CHA: Public history and its discontents or history in the age of Wikipedia. *Journal of the Canadian Historical Association/Revue de la Soci  t   historique du Canada* 18(1):1–26.
- Furet, F. 1971. Quantitative history. *Daedalus* 151–167.
- Geyer, M., and Bright, C. 1995. World history in a global age. *The American Historical Review* 100(4):1034–1060.

- Gieck, R.; Kinnunen, H.-M.; Li, Y.; Moghaddam, M.; Pradel, F.; Gloor, P. A.; Paasivaara, M.; and Zylka, M. P. 2016. Cultural differences in the understanding of history on Wikipedia. In *Designing Networks for Innovation and Improvisation*. Springer. 3–12.
- Hale, S. A. 2014. Multilinguals and Wikipedia editing. In *Proceedings of the 2014 ACM Conference on Web Science*, 99–108. ACM.
- Jackson, C. 2016. Using social network analysis to reveal unseen relationships in Medieval Scotland. *Digital Scholarship in the Humanities* fqv070.
- Jensen, R. 2012. Military history on the electronic frontier: Wikipedia fights the War of 1812. *The Journal of Military History* 76(4):523–556.
- Karimi, F.; Bohlin, L.; Samoilenko, A.; Rosvall, M.; and Lancichinetti, A. 2015. Mapping bilateral information interests using the activity of Wikipedia editors. *Palgrave Communications* 1.
- Kiser, E., and Hechter, M. 1991. The role of general theory in comparative-historical sociology. *American Journal of Sociology* 1–30.
- Lin, J. 1991. Divergence measures based on the Shannon entropy. *IEEE Transactions on Information theory* 37(1):145–151.
- Liu, J. H.; Goldstein-Hawes, R.; Hilton, D.; Huang, L.-L.; Gastardo-Conaco, C.; Dresler-Hawke, E.; Pittolo, F.; Hong, Y.-Y.; Ward, C.; Abraham, S.; et al. 2005. Social representations of events and people in world history across 12 cultures. *Journal of Cross-Cultural Psychology* 36(2):171–191.
- Luyt, B. 2011. The nature of historical representation on Wikipedia: Dominant or alternative historiography? *Journal of the American Society for Information Science and Technology* 62(6):1058–1065.
- Malkov, A. S. 2014. The silk roads: a mathematical model. *Cliodynamics: The Journal of Quantitative History and Cultural Evolution* 5(1).
- Michel, J.-B.; Shen, Y. K.; Aiden, A. P.; Veres, A.; Gray, M. K.; Pickett, J. P.; Hoiberg, D.; Clancy, D.; Norvig, P.; Orwant, J.; Pinker, S.; Nowak, M. A.; and Aiden, E. L. 2011. Quantitative analysis of culture using millions of digitized books. *Science* 331(6014):176–182.
- Müllner, D. 2011. Modern hierarchical, agglomerative clustering algorithms. *arXiv preprint arXiv:1109.2378*.
- Padgett, J. F., and Ansell, C. K. 1993. Robust action and the rise of the Medici, 1400–1434. *American Journal of Sociology* 1259–1319.
- Pentzold, C. 2009. Fixing the floating gap: The online encyclopaedia Wikipedia as a global memory place. *Memory Studies* 2(2):255–272.
- Pfister, D. S. 2011. Networked expertise in the era of many-to-many communication: On Wikipedia and invention. *Social Epistemology* 25(3):217–231.
- Potthast, M.; Stein, B.; and Gerling, R. 2008. Automatic vandalism detection in Wikipedia. In *European Conference on Information Retrieval*, 663–668. Springer.
- Rinser, D.; Lange, D.; and Naumann, F. 2013. Cross-lingual entity matching and infobox alignment in Wikipedia. *Information Systems* 38(6):887–907.
- Rosenzweig, R., and Thelen, D. P. 1998. *The presence of the past: Popular uses of history in American life*, volume 2. Columbia University Press.
- Rosenzweig, R. 2006. Can history be open source? Wikipedia and the future of the past. *The Journal of American History* 93(1):117–146.
- Rovira, D. P.; Deschamps, J.-C.; and Pennebaker, J. W. 2006. The social psychology of history: Defining the most important events of the last 10, 100, and 1000 years. *Psicología Política* (32):15–32.
- Rüsen, J. 1996. Some theoretical approaches to intercultural comparative historiography. *History and Theory* 5–22.
- Samoilenko, A.; Karimi, F.; Edler, D.; Kunegis, J.; and Strohmaier, M. 2016. Linguistic neighbourhoods: explaining cultural borders on Wikipedia through multilingual co-editing activity. *EPJ Data Science* 5(1):1.
- Samoilenko, A. 2017. Multilingual historical narratives on Wikipedia. Dataset. Version 1. GESIS Data Archive. <http://dx.doi.org/10.7802/1411>.
- Schich, M.; Song, C.; Ahn, Y.-Y.; Mirsky, A.; Martino, M.; Barabási, A.-L.; and Helbing, D. 2014. A network framework of cultural history. *Science* 345(6196):558–562.
- Sindbæk, S. M. 2007. The small world of the Vikings: networks in early medieval communication and exchange. *Norwegian Archaeological Review* 40(1):59–74.
- Spoerri, A. 2007. What is popular on Wikipedia and why? *First Monday* 12(4).
- Turchin, P.; Currie, T. E.; Turner, E. A.; and Gavrillets, S. 2013. War, space, and the evolution of Old World complex societies. *Proceedings of the National Academy of Sciences* 110(41):16384–16389.
- Turchin, P. 2011. Toward cliodynamics—an analytical, predictive science of history. *Cliodynamics* 2(1).
- Wertsch, J. V. 2002. *Voices of collective remembering: Test*. Cambridge University Press.
- Yu, A. Z.; Ronen, S.; Hu, K.; Lu, T.; and Hidalgo, C. A. 2016. Pantheon 1.0, a manually verified dataset of globally famous biographies. *Scientific data* 3.