Converging ‘Truths’?
The Role of Digital Knowledge Platforms for Information Flows across Language Barriers and History

Discord, polarization, and “parallel truths” on digital platforms are widely believed to threaten social cohesion, with serious repercussions for political and economic stability. The erosion of mutually agreed ‘truth’ on social networks risks undermining the foundations of societal consensus in democratic societies, prompting substantial empirical research in Economics and IS on polarization and filter bubbles. However, competing and contested narratives about past events are already familiar to Historians and Memory Studies, though these fields have yet to study narrative divergence quantitatively or at large scale.

We propose a groundbreaking, scalable quantitative analysis of how diverging versions of ‘the truth’ evolve and reconcile on digital platforms across different languages. By combining cross-lingual Large Language Models (LLMs), Artificial Intelligence (AI), modern econometric methods, scraping technologies, and API access to Wikipedia’s unique features as a language-specific knowledge repository, this approach enables unprecedented precision and scale in studying these questions.

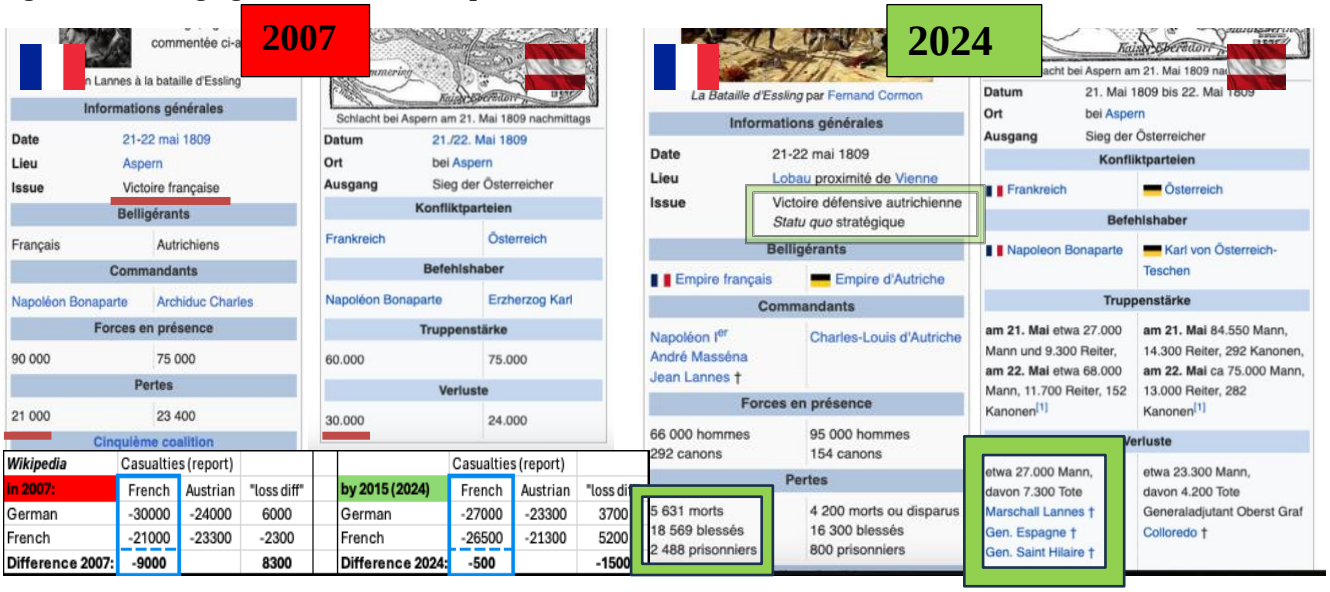
Specifically we develop a new methodology to contrast the monthly revision histories of over 1000 historic encyclopaedic entries (“topics”) in over 50 languages and for over 20 years (over 10,000,000 textual observations), to (1) quantify conflicts in knowledge, (2) analyse how these conflicts grow or diminish over time, and (3) shed light on factors and mediators which affect the conflict between and co-evolution of ‘truths’ across different languages.

Discord, Polarization and "parallel truths" on Digital Platforms are widely believed to threaten social cohesion, with potentially serious repercussions on political and economic stability in democratic societies.

Because the erosion of mutually agreed 'truth' on social networks could threaten the foundation for societal consensus, this phenomenon has inspired a sizeable body of empirical research about polarization on social media platforms in economics and information systems among other disciplines (Álvarez et al., 2020; Barberá, 2020; Choi et al., 2020; Cinelli et al., 2021; Garimella et al., 2018; Greenstein & Zhu, 2012, Greenstein & Zhu, 2018; Kitchens, Johnson, and Gray, 2020; Terren and Borge-Bravo, 2021). However, the idea of competing and contested narratives about past events is highly familiar in History and Memory Studies, where cultural memory and historiography as such are construed as fluid and selective processes that are subject to media as well as societal and cultural factors (Halbwachs, 1992; Erll, 2011; Olick, 2011; Assmann, 2011; Garde-Hansen, 2011; Ricoeur, 2004). Conflicts, wars, and pivotal events are understood to be documented from perspectives shaped by nationalistic, cultural, or political frameworks (e.g., Leggewie and Lang, 2011).

Simple examples include differing names for events ("Fall" vs. "Conquest of Constantinople"; "Reconquista" vs. "Fall of Al-Andaluz") and battles with both sides using a different field names, casualty numbers or even both sides claiming victory ("Schlacht bei Aspern" vs. "Bataille d'Essling"). Older French sources report 21,000 French casualties at Aspern-Essling, while Austrian sources state 30,000 French casualties. Similarly, for the Battle of Trafalgar, French accounts cite 26,000 French vs. 18,500 English troops (ratio: 71%), while English sources claim 30,000 French vs. 17,000 English (ratio 54%).

Figure 1: Converging Narratives on Wikipedia: Battle Boxes



Wikipedia	Casualties (report)				Casualties (report)		
in 2007:	French	Austrian	"loss diff"	by 2015 (2024)	French	Austrian	"loss di
German	-30000	-24000	6000	German	-27000	-23300	3700
French	-21000	-23300	-2300	French	-26500	-21300	5200
Difference 2007:	-9000		8300	Difference 2024:	-500		-1500

Notes: *Figure 1 contrasts the French and the German "BattleBox"-summaries for the Battle of Aspern-Essling. The lef panel shows the information presented in 2007, the right side corresponds to December 2024. The box at the bottom summarizes the respective casualty numbers provided by each side at both points in time. While the divergence in the loss difference is 8300 in 2007, this is reduced to 1500 in 2024 (and by 2015). Divergence on French casualties is reduced from 9000 to 500.*

In this proposal, we will refer to this phenomenon as "***divergence in 'truths'****,"* because both accounts typically reflect the language-specific tradition of historiography and can be traced back to reputable encyclopaedic sources. Numerous other examples exist of different languages keeping diverging records of historic events of conflict, such as battles and other atrocities (e.g., Leggewie and Lang, 2011). but this divergence has – as of yet - not been studied quantitatively and at large scale. Yet, such biases, deeply embedded in historical accounts foster misunderstanding and could hinder both reconciliation and even economic development.

An example of harmonization or converging 'truths' are the German and French Wikipedia versions of the battle of Aspern-Essling. Austrian history remembers "Schlacht bei Aspern" (e.g.: 'Lion of Aspern' in Vienna), while in French history remembers "Bataille d'Essling" (e.g., 'Arc de Triomphe' in Paris). Initially, neither Wikipedia page mentioned the other name, but now both do. Until 2007, "Bataille d'Essling" stated a French victory, while "Schlacht bei Aspern" stated an Austrian victory. Since 2007, "Bataille d'Essling" describes it as an "Austrian victory without strategic consequences." The divergence in French casualties has reduced from 9000 to 500 (now: "Bataille d'Essling": 26,500, "Schlacht bei Aspern": 27,000).

While the existing literature has focused on the potential of digital platforms to amplify discord and increase conflict, the present research proposal posits that connecting and juxtaposing conflicting historic narratives can foster a process of harmonization and contribute to a resolution and convergence of long-standing conflicting narratives. We thus depart from the existing "one-way paradigm" of amplified conflict to analyse whether digital platforms also offer novel opportunities to enable a mechanism of reconciliation. However, we stress that the hypothesized mechanism is fundamentally different from solutions suggested in the existing literature on echo chambers (e.g. Terren and Borge-Bravo, 2021).

To address this research challenge, we propose to focus on the Wikipedia ecosystem which offers a unique laboratory because every edit is recorded, a feature used extensively in the PIs own past research contributions (Halatchliyski et al. (2016); Kummer et al. (2021), Hinnosaar et al. (2022, 2023)). Moreover, Wikipedia provides a platform that allows for language-specific divergence in narratives and provides us with a research context of clearly connected (by topic and language links) yet unambiguously separated narratives (across languages), that is perfectly suited for developing the proposed methodology. In that sense, Wikipedia offers a unique laboratory to study and quantify both of the following phenomena: (i) the existence, transmission, and propagation of language-specific narratives as an instance of "historic echo chambers" and (ii) the extent to which digital knowledge platforms can (or fail to) induce a process of convergence in narratives or "converging 'truths'." To tackle these two questions, we propose to leverage multilingual Large Language Models (LLM, aka 'AI'), and access to Wikipedia's multilingual revision histories to analyse contested content, such as

monthly accounts of over 500 past battles and 500 neutral topics in over 50 languages. Thus, obtaining a dataset of over 10 million observations (and potentially more), we can explore whether digital platforms can facilitate a process of convergence for pre-existing language-specific differences in historical 'truths,' before extending the methodology to other domains.

Literature Review

The potential for bias in online platforms is well-documented. Prior literature highlights in-group bias across languages (Álvarez et al., 2020) and political slants away from neutrality in both expert-curated Encyclopaedia Britannica and crowd-sourced Wikipedia (Greenstein & Zhu, 2012, 2018). We conceptualize **historic biases and diverging accounts of historical events** between languages and nations as a "**historic echo chamber effect**," where cultural perspectives and preferences amplify differences in narratives.

The echo chamber effect has been extensively explored in social media, revealing mechanisms such as algorithmic filtering, self-selection, and homophily that reinforce pre-existing beliefs (Cinelli et al., 2021). Studies further show how social media induces partisan shifts, amplifies rumors, and reduces bipartisanship (Kitchens et al., 2020; Choi et al., 2020; Barberá, 2020; Garimella et al., 2018). These dynamics contribute to polarization and challenge democratic discourse, while mitigation strategies, such as diversifying information sources, have been explored (Terren & Borge-Bravo, 2021). These findings ***resonate with cornerstone insights from*** the disciplines of ***history*** and ***memory studies*** who highlight that memory, the recollection of facts, and historiography as such are fluid and selective processes that are subject to societal and cultural factors (Halbwachs, 1992; Erll, 2011; Olick, 2011.), which interact in complicated ways with the medium that creates the record. (Assmann, 2011; Garde-Hansen, 2011; Ricoeur, 2004).

Theories of truth—coherence, consensus, and pragmatic—offer frameworks for analyzing information convergence. Coherence theory emphasizes logical consistency within interconnected beliefs (Rescher, 1973; Kirkham, 2001), while consensus theory defines truth through rational agreement in a community (Habermas, 1984; Rescher, 1993). Pragmatic theory ties truth to practical consequences, viewing it as utility-driven and evolving with new information (James, 1907; Peirce, 1878). Together, these theories inform how data from multiple sources can merge into a unified narrative (Blanshard, 1939; Jenkins, 2006). The present research proposal posits that ***connecting*** and ***juxtaposing*** conflicting historic narratives ***can foster a process of harmonization*** and contribute to a ***resolution of conflict***. We argue that digital platforms offer a novel opportunity for bringing this mechanism to fruition. However, we stress that the hypothesized mechanism is fundamentally different from solutions suggested in the existing literature on echo chambers (e.g. Terren and Borge-Bravo, 2021). ***Wikipedia offers a unique laboratory*** to study and quantify both phenomena: (i) the existence, transmission, and propagation of language-specific narratives as an instance of "historic

echo chambers" and (ii) the extent to which digital knowledge platforms can (or fail to) induce a process of convergence in narratives or "**converging truths**."

The value of understanding this channel cannot be overemphasized, as historical grievances and differing interpretations of history have been sources of tension and conflict between ethnic groups and nations (see e.g., Galtung, 1996).. In contrast, a shared vision of history can foster cooperation and diplomacy and promote stability and peace. (e.g., Larson & Shevchenko, 2010, Montville, 1990). This perspective has also been at the heart of famous efforts including the Gacaca courts in Ruanda (Clark, 2010; Minow, 1998) or truth and reconciliation commissions (Doxtader, 2009; Tutu, 1999; Wilson, 2001). Related research found that providing students with a cohesive understanding of their country's past promoted national unity, social cohesion, and a sense of identity (e.g., Noddings and Brooks, 2016), while shared historical narratives facilitate cultural exchange and tourism between nations. (e.g., Timothy & Boyd, 2006). We take this as overwhelming evidence for the value of promoting a shared narrative of historical accounts.

Wikipedia is the ideal laboratory to study whether digital knowledge platforms induce convergence in narratives, because Wikipedia constructs knowledge collaboratively, with consensus emerging through, coherence, and practical relevance (Kittur & Kraut, 2008; Jemielniak, 2014; Reagle, 2010). Talk pages foster rational dialogue, aligning with consensus theory (Habermas, 1984), while coherence is ensured by integrating new information with existing content (Forte & Bruckman, 2008). Pragmatic utility guides contributions to prioritize relevance (Fallis, 2008). Mechanisms like Creative Commons licenses (Lih, 2009), verifiability policies (Magnus, 2009), and collective editing on talk pages balance differing viewpoints to create unified, unbiased narratives (Forte et al., 2012; Reagle, 2010).

Moreover, Wikipedia's multilingual structure uniquely reflects diverse cultural perspectives, with independently operated language editions often differing in scope and historical interpretations (Hecht & Gergle, 2010). Tools like interlanguage links, multilingual editors, and translation aids, supported by Wikidata, help bridge gaps, streamline updates, and enhance synchronization while highlighting cultural differences (Adafre & de Rijke, 2006; Miquel-Ribé & Laniado, 2018; Vrandecic, 2012; Samoilenko et al., 2017). On the other hand, research has shown that the process is subject to challenges, including resource disparities, cultural biases, and misinformation. Smaller editions lag in updates (Lewoniewski et al., 2017), while edits may reflect national pride or differing sources (Pfeil et al., 2006). Wikipedia's inclusivity also invites unreliable edits, requiring strong editorial practices to ensure verifiability and consistency (Ford et al., 2013; Luyt & Tan, 2010). Language links aid synchronization and information transfer, but full convergence remains limited by resource and contextual differences (Adar et al., 2009; Callahan & Herring, 2011; Liu et al., 2018).

To summarize, our work builds on and contributes to four connected streams of literature:

- **Digital Platforms, Echo Chambers and Bias**: Studies highlight the prevalence of biases in digital platforms like social media platforms, and Wikipedia, particularly in early article versions, with later revisions often reflecting more balanced perspectives (Greenstein & Zhu, 2018; Kim et al., 2016). However, systemic biases remain, shaped by editorial practices and cultural perspectives. **Contribution of this project: *Truthconverge* breaks new ground** by proposing to analyse this question across many languages and at large scale.

- **Historiography and Narrative Bias**: Historical accounts, especially those documenting conflicts, frequently reflect the perspectives of victors or dominant cultural groups (Bloch, 1954; Carr, 1961; Tosh, 2015). This reinforces inequalities in representation and complicates reconciliation efforts. (Halbwachs, 1992; Hunt, 1989; Minow, 1998; Tutu, 1999) **Contribution of this project**: ***Truthconverge*** is groundbreaking, because we analyse the evolution of diverging narratives *over time* for a large body of conflicts simultaneously comparing *dozens* of languages.

- **Peace-building and Shared Histories**: Research which underscores the importance of establishing shared historical narratives in fostering reconciliation and empathy (Bekerman & Zembylas, 2011; Galtung, 1996). Digital platforms represent a novel avenue for promoting such shared understandings, and thus contribute to peace building and collaboration between language groups and societies. (Clark, 2010; Minow, 1998; Tutu, 1999; Zehr, 2002). **Contribution of this project**: ***Truthconverge*** is first to analyze this potential at large scale and to investigate whether very long standing discrepancies (100s of years) can be overcome by digital platforms..

- **Recent progress in Large Language Modelling**, has fueled AI-supported translation and Natural Language Processing methods. Major progress in NLP was achieved by several algorithms, including word2vec (Mikolov et al., 2013), BERT (Devlin et al., 2018), and – most recently -Multilingual Embeddings (Cohere 2024). Cohere, founded by former Google Brain employees in 2019, focuses on LLMs enabling access to researchers through its nonprofit research lab and open-source community, Cohere for AI. **Contribution of this project: *Truthconverge*** breaks new ground by systematically applying such models to study the problem of narrative convergence on topics of historic conflict.

Motivational Example: Battle near Aspern-Essling (1809)

In 1809 a battle between Austrian and French forces took place between villages called Aspern and Essling (now part of Vienna). While the battle saw no decisive victory Austrian forces managed to deny Napoleon's army a crossing of the danube river. After heavy losses on both sides, Napoleon's army retreated to the south, which was achieved without major defeat thanks to a strategic stronghold in Essling.

Convergence in overall narrative – Dependent Variable 1. The battle's diverging names—"Schlacht bei Aspern" in Austrian historiography and "Bataille d'Essling" in French accounts

highlight longstanding offline differences, which are reflected in history books and monuments, such as the Arc de Triomphe in Paris. These discrepancies persisted in early Wikipedia versions, with the French page recognizing the German name in 2004, but claiming a "victoire francaise" until 2007, before revising the battle outcome several times. The German page, conversely, emphasized Napoleon's defeat from the start and only mentioned the French name in 2016. Multilingual text embeddings and cosine similarity measures will quantify these narrative shifts.

Convergence in reports of Casualties - Dependent Variable 2: To illustrate the measure of convergence as change in the difference in casualties, we revisit casualty-reports about the Battle of Aspern-Essling in Figure 1. In 2007, German pages reported 24,000 Austrian vs. 30,000 French casualties (difference: -6,000), later revised to 23,300 vs. 27,000 (-3,700). French pages intermittently reported 21,000 French vs. 23,300 Austrian casualties (difference: -2,300) but by 2014 revised to 26,500 French vs. 21,300 Austrian casualties (+5,200). These patterns, including the reversal of reported differences, will be used to quantify narrative convergence. Irrespective of the "objective truth," the divergence in counts diminished from 8300 to 1500, an instance of "converging 'truths'." Moreover, the French 'truth' has changed more than the German 'truth.'

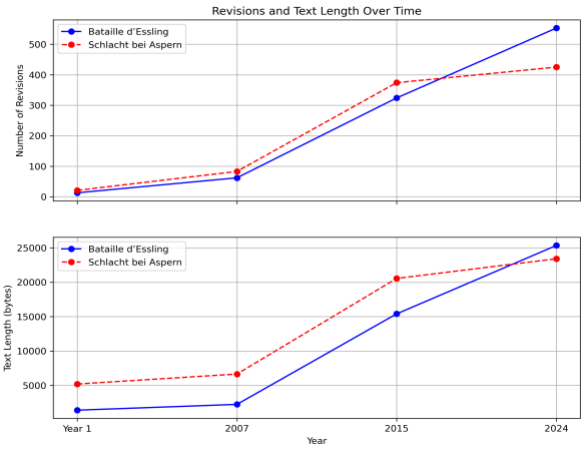


Figure 2: Content Availability and Edit Density

Text Density – Dependent Variable 3: Similarly, the density of editing and available information varies systematically by party. Winning parties tend to remember better and provide more content about event at hand than the losing parties. Looking at the Battle of Aspern-Essling again, we see that until 2015 both amount of content and number of edits were substantially higher on the German version of the article, which - especially until 2007 – featured three times as much text as the French version. Bearing in mind that information availability is endogenous, as involved parties might choose to discourage memory about unfavorable (for governing parties) events or conflicts. To account for this endogeneity, we will use measures that capture differences in the amount of available information as a third group of outcome variables to consider. The simplest of these measures can be obtained by computing the length difference or the ratio that results from comparing the article of the winning and the losing language.

Three stylized facts: We identify three key stylized facts that will guide our analysis. First, conflicting parties report events differently, often highlighting their own achievements while downplaying those of their adversaries. Second, the winning side tends to provide more detailed accounts, leveraging its position to shape the narrative—though this is less evident in cases involving crimes against humanity. Finally, discrepancies in accounts, including numerical data such as casualty

figures, are common, reflecting the subjective nature of historical reporting. These discrepancies often stem from biases, unequal access to information, or differing political and cultural agendas. For example, casualty figures may vary depending on whether reported by the victor, aiming to downplay losses, or the defeated, highlighting the damage inflicted.

Before moving to the general strategy of the analysis, we stress that the Example from the Battle of Aspern-Essling is not unique. A similar pattern (though under opposite signs) emerges for the Battle of Wagram which ensued only 6 weeks later. Other prominent examples are the "Fall"/"Conquest" of Constantinople and of Al Andaluz/Spain, or the Battle of Trafalgar and many more. It is the research goal of ***TruthConverge*** to analyze to which extent we can highlight this pattern to be systematic. The second major goal is to shed light on the relative positioning of "winning languages", "losing langauges" and uninvolved languages.

Figure 3: "Conquest"/"Fall" of Constantinople in Turkish, Greek/Romanian, and English 'Wikipedia eyes'



NOTES: *Main images chosen for the "Conquest" (turkish) or "Fall (Greek/Romanian/English) of Constantinople (1453) in different Wikipedias. While Turkish Wikipedia focuses on the emperor, the Image on Greek Wikipedia emphasizes the fortifications and bloodiness of the battle and emphasizes religion. English Wikipedia shows an image that emphasizes the strategic situation and important details of the naval battle.*

Main Questions: Our analysis is shaped by three guiding questions, emerging from the examples presented above. First, (1) do biases in reporting grow or diminish over time? Understanding this dynamic can reveal whether historical narratives tend to stabilize or diverge over time. Second, (2) does interlinking on platforms foster a "convergence of truths," where conflicting perspectives align through shared access and connectivity? Third, (3) do these biases, deeply rooted in cultural and historical contexts, persist and transmit through centuries, shaping how societies remember and interpret the past? In addition, we ask: what is the role of "neutral" languages in bridging these divides and mediating between conflicting narratives?

RESEARCH STRATEGY

To address this set of questions we propose to organize the research in five work packages that mutually complement each other to form a balanced overall project. ***Subproject 1 corresponds to Aim 1*** and provides the fundamental measurement technique and quantification that will be used in subsequent Work Packages. ***Subproject 2 (Aim 2): Analysis of 'baseline convergence' and platform interventions*** will deploy the methodology developed in subproject 1 to (1) quantify and analyse

knowledge conflicts across languages, and (2) to analyse the effect of platform strategies such as cross-language linking or AI-driven translation tools on narrative convergence. ***Subproject 3 (Aim 3): Investigate root causes and exogenous drivers of Converging 'Truths.'a***ims to analyse driving factors—cultural, technological, and editorial—that promote or hinder narrative alignment across languages. ***Subproject 4 (Aim 4): Agency and micro-dynamics of editing behaviors,*** will focus on the endogenous micro-dynamics of editing behaviors to understand individual and collective agency in fostering (or impeding) converging 'truths,' In the final ***Subproject 5 (Aim 5): Validation and Application to single-language contexts,*** wil deploy a "sanity check" and validate the insights from this and the other subprojects (1, 3 and 4). A field-test interventions aims to roll out a warning for language versions that significantly depart from the majority consensus.

In what follows, we will outline each of the five subprojects and the data collection procedure in greater detail.

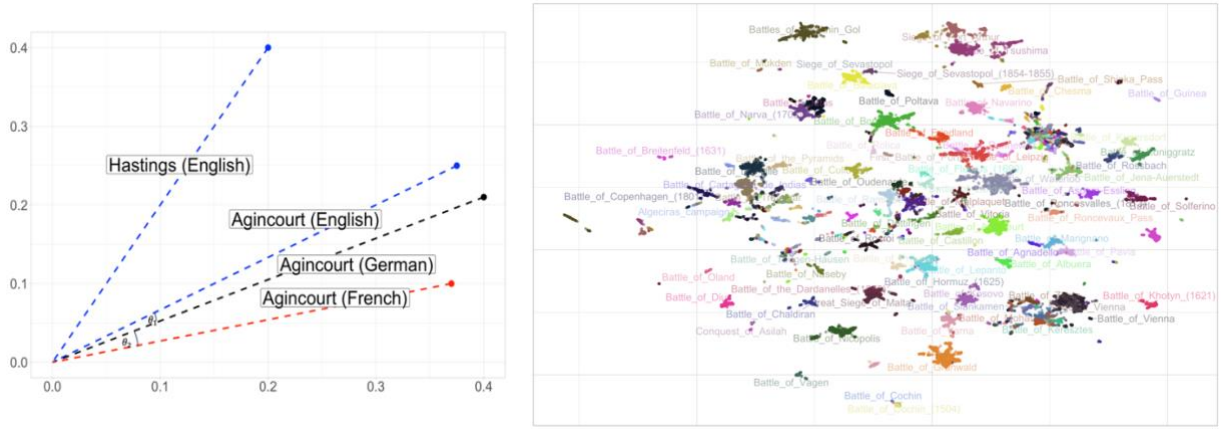
Subproject 1 (Aim 1): Quantifying Language-Specific Narratives and Converging 'Truths':

Subproject 1 forms the backbone of the entire research project. Its overall aim is to develop a universal and scalable quantification of discrepancies in how different languages document historical conflicts and to identify patterns and trends over time. This package will leverage recent NLP methods, multilingual LLM embeddings, scraping and parsing techniques, sentiment analysis, and computational similarity measures. Embedding-based similarity measures will provide a key metrics for measuring convergence in narratives.

Multi-Lingual LLM Embeddings and Cosine Similarity:

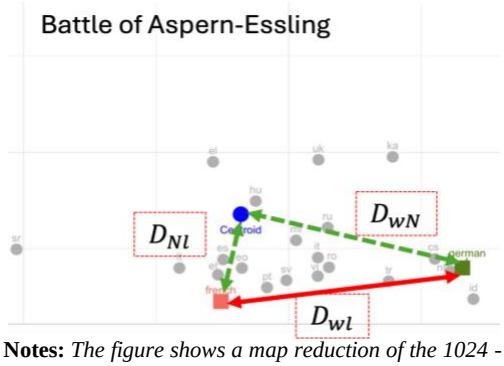
The core idea of ***TruthConverge*** is to analyze how the semantic differences in historical narratives across languages evolves over time. To achieve this goal, ***TruthConverge*** leverages **Natural Language Processing (NLP)** techniques and **multi-lingual large language model (LLM) embeddings.** The core idea is to embed all language-versions of an article in one multi-lingual semantic space and to quantify the alignment or divergence of narratives. The quantification is achieved by evaluating their **cosine similarity**, a measure of how closely two text embeddings are aligned in a high-dimensional space

Figure 4: Illustration of cosine similarity and initial mapping of cross-lingual similarity scores (of 2023 Wikipedia version)



This is illustrated in the left panel of Figure 4 with a visual representation based on two historical battles: the **Battle of Hastings** and the **Battle of Agincourt**. The figure contrasts **semantic cross-language distance in narratives** with same-language semantic distance between two topics (e.g. different battles). The English account of the Battle of *Hastings* shows significant semantic distance to all narratives about the Battle of *Agincourt*, as reflected in its isolated position on the graph. Among the more closely aligned multi-lingual accounts of the Battle of Agincourt the English and German versions show closer alignment, while the French version diverges more significantly. This illustrates how the proposed methodology will allow to quantify how historical events are framed differently across languages, shaped by cultural or lingusitic context. For example, the French narrative of Agincourt may emphasize aspects that differ from those in the German or English accounts, reflecting the perspectives of each nation's historiography.

Figure 5: Variable 1: distance of narratives.



Notes: *The figure shows a map reduction of the 1024 - dimensional cosine similarity measures for Battle of Aspern-Essling in 2023. The squares show the language versions of the conflicting parties (German and French), the grey dots represent articles in neutral languages. The "Centroid" is the average location of neutral languages. The three arrows* D_{wl}, D_{wN}, D_{Nl} *indicate the distance measures that will be used as main*

Initial Cross-Language Comparison via LLMs (2023 versions) : The left panel of Figure 4 presents the results of embedding historical battles into **multi-lingual large language model (LLM)** by **Cohere** and visualizing the relationships using dimensionality reduction techniques. The analysis explores how historical narratives differ across languages by mapping the cosine similarities and differences of battle descriptions. Colors represent battles (e.g., **Battle of Agincourt**, **Battle of Vienna**) and points represent Wikipedia pages in a given language (**German, French, English, Swahili, etc.**) as embedded by Cohere's multi-lingual LLM. The spatial distribution reflects cosine

similarities. Closer points indicate greater alignment in narrative framing, while more distant points suggest divergence.

Clusters indicate events described with similar semantic framing across languages. For instance, tightly grouped battles suggest a shared or convergent narrative across linguistic boundaries, while dispersed points highlight divergence in how different languages frame the same event. Some battles, like **Battle of Vienna**, appear closely aligned across languages, reflecting greater semantic agreement.Others, like **Battle of Tsushima** or the **Siege of Port Arthur**, are more dispersed, indicating greater variation in how the event is described in different languages.

Figure 5 shows the cluster corresponding to the Battle of Aspern-Essling. A triangular distance measure between averages of winners, losers, and neutral languages provides key metric for analysing narrative (static) *distance.* (cf. part B2). The red-blue triangle shows the quantification of the three distances between winning party, losing party and the average uninvolved language.

Analysis of convergence: This measure evaluating narrative distance. We start with how winners diverge from losers at the article level (D_{Awlt}) at the beginning and then estimate linear and flexible time-trends, to evaluate whether divergences grow or diminish **over time**. We contrast the battle specific distance with analogous distance measures for the articles about planets and elements between the languages *w* and *l.* A simple specification to estimate group specific linear trends at the topic level *A* is given by:

$$D_{Awlt} = \beta_0 + \beta_1 Battle_{Awl} + \beta_2 t + \beta_3 Battle_{Awl} \times t + \varepsilon_{At} \quad (1)$$

D_{Awlt} is the measure of narrative distance between the article *A* in language of the winner *w* and the loser *l*. $Battle_{Awl}$ takes the value 1 if the article is about a battle and if it was fought between *w* and *l*, and t is a simple numeric variable representing time. Suitable non-parametric alternatives will be estimated using time as a categorial variable and estimating non-parametric methods including splines or kernel smoothing.

The analysis of D_{Awlt} will be underpinned by an analysis of the alternative distance measures that are given by the differences in article properties (e.g. length) and casualty reports (compared against numbers in the info boxes of the control group). In addition to the control groups based on planets and elements, we can also use battle pages for the same languages but about other battles that did not involve *w* and *l* as control groups.

In step 2, we can leverage the availability of uninvolved ("neutral") languages. Using the centroid (i.e. average) of the measure obtained for these languages and repeat the analysis but contrast the winner's distance to the centroid, D_{AwNt}. to the loser's distance D_{AlNt}. Doing so allows to analyse initial distances and convergence between each party and the position of the average uninvolved

party. Moreover, the contrast of the convergence between neutral and winner vs. convergence between neutral and loser will shed additional light on the convergence results (or absence thereof) in D_{Awlt}. A third aggregate measure is given by the circumference of the triangle $D_{Awlt} + D_{AwNt} + D_{AlNt}$, which can be conveniently tracked over time.

Using multi-lingual LLM embeddings, like those from Cohere, combined with Wikipedia's revision history to track changes at a monthly level, represents a **groundbreaking quantitative and scalable method** for studying narrative convergence and divergence across languages over time. By moving beyond surface-level differences in wording to capture deeper shifts in meaning and the evolving convergence of **'truths,'** this method provides **unparalleled insights** into **how digital platforms can foster convergence in long-standing, diverging historical narratives.** Unlike the *familiar* discord and divergence often seen on social media, this approach explores how such platforms can **constructively reshape collective understandings of the past** and how **language, culture, and economic factors** influence this process.

Subproject 2 (Aim 2): Analysis of 'baseline convergence' and platform interventions

In subproject 2, we will deploy the methodology from subproject 1 to (1) quantify and analyse knowledge conflicts across languages, but, more importantly (2) to analyze the efficacy of platform strategies such as cross-language linking or AI-driven translation tools in fostering narrative convergence. As before, the initial package focuses on the analysis of historic accounts of the same battle across different languages and whether initial discrepancies in the different accounts on Wikipedia diminish over time (or do not) because of – among other things – exposure to the other language versions about the same battle.

We will leverage three major design features that were introduced by the platform since 2002 and test how (i) interlinking, (ii) user interface standardization (battleboxes), and (iii) AI-driven tools for translation affect biases and foster alignment. For all three interventions, this part will deploy event study techniques, and thus a short description of the interventions is in order.

The **first intervention**, "battle boxes," are standardized templates summarizing battles in the top right corner of articles (e.g., "Victoire francaise" vs. "Austrian victory" or casualty discrepancies). These templates make stark discrepancies more visible, potentially triggering convergence. Battle boxes are user-added templates, enabling two event study moments: (1) **Template Availability**: When the template becomes available across a language, allowing estimation of an ITT (Intention to Treat) effect. (2) **Template Activation**: A battle box is added to a specific article, allowing ATT (Average Treatment Effect on the Treated) estimation. Identification arises as different language versions (e.g., Dutch or Swedish) activate the battle box at different times, enabling comparisons across language-battle observations. In addition, uncontested content, such as geometric shapes or celestial bodies, serves as alternative comparison group.

The **second intervention**, **language links**, connects articles to their counterparts in other languages, making cross-language versions one click away. Like battle boxes, language links make discrepancies more visible, triggering convergence. This intervention also supports ITT estimation (when the feature is generally available) and ATT estimation (when links are added). As before, control observations are provided by other language versions and by control topics. The **third intervention**, **AI translations**, is rolled out universally across the platform, making treatment exogenous. However, unlike the first two interventions, other languages cannot serve as controls. Instead, we use uncontested content, such as geometric shapes or celestial bodies, as a comparison group to control for baseline differences between languages. This comparison group is also applicable to the first two interventions.

Estimation strategy: Each intervention will be studied independently, but the estimation methodology is consistent across them. We define "TreatAct" as an indicator variable equal to 1 when a feature (e.g., Battleboxes) is added to an article or article pair, signifying active treatment. While features can be removed, they are typically improved over time.We hypothesize that treatment accelerates convergence, with the null hypothesis being no effect on narrative convergence. To test this, we estimate a DiD regression using article distance measures (e.g., article length, cosine similarity, or – for battles – info-box numbers). For simplicity, we focus on battle analysis using planets as a control group and Language Links as treatment. The methodology generalizes to other topics like historical figures. Our sample includes all battle-language and planet-language observations, forming a topic-language-month panel where an article A in language L is the observational unit. For article length, we construct a balanced monthly panel, coding non-existent articles as zero. For other outcomes, the panel is balanced for existing articles (cosine similarity) or reported casualties (casualty differentials). The simplest two-way fixed-effects:

$$\Delta D_{Awlt} = \beta_0 + \beta_1 TreatActive_{Awlt} \times Battle_{Awl} + \delta_A + \gamma_t + \varepsilon_{ALt} \quad (2)$$

This regression can be expanded to a full event study design with period-battle interactions and time normalized to 0 at the time when treatment between *w* and *l* is activated. D_{Awlt} is the measure of narrative distance between the winner's (*w*) and the loser's (*l*) version of article *A* and the centroid of the neutral languages in month *t*. For the event-study analysis we focus on the change in narrative divergence by taking the first difference to the month before. $TreatActive_{ALt}$ equals 1 if and only if treatment is active for the given article *A* at time t (e.g. if a BattleBox is present or a language link is in place) and, otherwise, equals 0. $Battle_{Awl}$ is an indicator variable taking the value 1 if the article is a battle-observation involving languages *w* and *l* as conflicting parties, and 0 if it is from the control group. Article δ_A and weekly fixed effects γ_t are also included. As before, we can expand to compare the winner's and loser's distance to the neutral language centroid (N_A) for article *A* (D_{AwN_At} *vs* D_{AlN_At}) and include article-language level fixed effects α_{AL} for the winner and the loser

respectively (not shown here).

While each event study is an independent project, this subproject contributes to the literature on information systems, because each intervention uses a different channel to make conflicting information more salient. Language links improve access to other versions, battle boxes highlight key contradictions (e.g., in casualties), and AI translations expand the pool of users accessing foreign-language information. Comparing how the effects differ for each channel offers insights into their relative importance in exposing conflicting narratives.

Specification issues: A key challenge in our analysis is the potential endogeneity of language links and battle boxes, as they are placed by users. First, reverse causality may occur if a link is added due to prior convergence. To address this, we use changes in narrative divergence as the dependent variable and run regressions with increasingly small time windows around the event, ensuring convergence is not misattributed. Pre-trend analyses further confirm that convergence does not precede the intervention. In SP 4, edits co-occurring with the placement of links will be studied in detail to understand the role of individual editors versus the crowd.

Second, the user who places a link may also edit articles, systematically reducing divergence. To account for this, we: (1) exclude convergence attributable to the user who placed the link, (2) exclude edits occurring within a week or month of the intervention and focus on changes in convergence speed beyond this window, and (3) conduct "honest DiD" analyses of pre-trends (Rambachan and Roth, 2019, 2023).

Lastly, battle boxes and language links are introduced at varying times, creating heterogeneous treatment timing. Standard Two-Way Fixed Effects (TWFE) estimators may produce biased results due to improper weighting of treatment effects, particularly when earlier-treated groups act as controls. To address this, we apply the Callaway and Sant'Anna (2021) methodology, which adjusts for heterogeneous treatment effects and avoids bias from incorrect control group weighting.

Feasibility and Risk Management Subproject 2

The risks associated with this study are low. The main concern would be that Wikipedia might go offline, block scraping or remove the revision history. To minimize the potential impact of any such changes we have been pre-scraping the information for the previously mentioned pilot study (at a quarterly level for 104 battles, 200 celestial bodies and in 12 languages) and we will continue to expand the amount of pre-scraped data, expanding the set of battles, languages and control pages.

Subproject 3 (Aim 3): Investigate root causes and exogenous drivers of Converging 'Truths.'

The goal of subproject 3 is to broaden the analysis and investigate driving factors—economic, technological, cultural and editorial—that promote or hinder narrative alignment across languages, proceeding in two parts. Part 1 **broadens the analysis to other contested topics** (e.g., rulers, places,

cultural objects) to explore domains and languages where narratives are more (or less) likely to converge. This step incorporates language-specific macroeconomic and societal variables—such as growth, trade, human development, press freedom, conflict recency, and trade relationships—to identify predictors of narrative convergence. Understanding these patterns is crucial for contextualizing narrative dynamics across domains and languages.

The implementation of the analysis is a straightforward extension of the models in Equation (1) and (2). For example, to analyse the moderating effect of trade-intensity, we can generate categorial variables for low and high intensity trade partners, and then use interaction terms to estimate group-specific heterogeneous effects.

Part 2 examines **sources of language-specific narrative divergence**, contrasting national historiographic traditions and early Wikipedia versions. By analyzing digitized traditional encyclopedias (e.g., Britannica, Larousse, Brockhaus) and comparing narrative distances between (a) traditional sources, (b) traditional sources and early Wikipedia versions, and (c) traditional sources and later Wikipedia versions, this part reveals whether Wikipedia convergence reflects shifts in narratives or correction of early biases. This approach links traditional sources to Wikipedia to assess the origins and persistence of historiographic biases on digital platforms.

After digitizing the original historic encyclopaedias, the methodology from SP1 can be applied directly to measure and analyse narrative distances between historic encyclopaedic, early and late Wikipedia.

Feasibility and Risk Management Subproject 3:

Part 1 does not imply any additional risk, once we can use the method developed in subproject 1, because the macroeconomic data and human development indices are widely available and can be easily merged. For part 2, the key challenge lies in identifying traditional paper encyclopaedias and other relevant traditional sources and obtaining digitized versions that can be used in the analysis. This task will be undertaken by the team and with support of library services at Nova and the advisory board members.

Subproject 4 (Aim 4): Agency and micro-dynamics of editing behaviors

This part focuses on the micro-dynamics of editing behaviors to understand individual and collective roles in fostering or impeding narrative convergence. It will also develop a measure for the "heat" or conflictiveness behind diverging narratives, enabling analysis of within-language contexts. Using edit-level information on Wikipedia, we analyze how individual edits and broader editing patterns affect convergence. By quantifying the narrative distance between revisions (e.g., Revision 208 vs. 209) and determining whether edits move the text closer to or further from other language versions, we can measure conflict in editing activity. For instance, dividing total absolute change by cumulative narrative change in an article identifies (i) editing conflicts with neutralizing changes or (ii) a

systematic push toward a different narrative. This enables exploration of what (or who) drives convergence and what impedes it, including the role of multilingual editors or "edit wars."

We address two key questions: Are convergence events (e.g., “Bataille d’Essling in 2007”) collaborative or contentious? Does non-convergence stem from inactivity, or from rejected attempts at harmonization?

Second, while Wikipedia’s separation of narratives across languages is ideal for our method, this subproject will generalize to contexts where conflicting groups within the same language edit the same article (e.g., Brazil vs. Portugal, Spain vs. South America, American vs. English, or Austrian vs. German editors). We will identify group-specific patterns (e.g., Brazilian, conservative, or religious editors) by linking editing behaviors to text bodies and keywords associated with these groups, following prior work (e.g., Álvarez et al., 2020; Greenstein & Zhu, 2018). Editors will be associated with groups similar to how they were linked to languages, enabling an analysis of their agency and narrative shifts within one language. This will serve as a first step toward understanding conflicting narratives and their convergence within same-language contexts.

Feasibility and Risk Management Subproject 4:

The core risk in this step lies in the necessity of profiling individual editors, which is circumvented by analysing monthly article-observations in SP1, SP2 and SP3. Even though authors are anonymized by usernames, and the profiling is provided by Wikipedia, we will take extra care to work with aggregated data whenever possible and avoid explicit username-profiling by employing an anonymous numerical, hexadecimal (or similar) individual key that disallows associating any user-observation in the dataset to a username on Wikipedia.

Subproject 5 (Aim 5): Validation and Application to single-language contexts

In this subproject we will (a) validate the research from earlier subprojects and (b) attempt to generate value to Wikipedia and society by taking the concepts and insights established in subprojects 1 to 4 to application.

Two methods will be deployed to “sanity check” and validate the insights from this and the other subprojects (1, 3 and 4). The first one is a field-test intervention with the aim of tackling strong language-specific biases in the Wikipedia narrative. To achieve this, we and my team will design and roll out a warning for language versions that significantly depart from the majority consensus in other languages and then analyse in a field experiment approach whether such an intervention influences readership and/or induces (an attempt at) convergence in the narrative. In addition to the field experiment, qualitative interviews with Wikipedia editors will be conducted in a structured survey. Both methods together will validate findings and predictions from the other work packages through real-world applications on Wikipedia and similar platforms.

Second, an exploratory study will further apply the insights from previous subprojects to other social media platforms. Ultimately – and with a high risk/high gain outlook - the tools developed here, may potentially be capable of detecting sudden changes in the dynamic of conflict, such as a growing likelihood of conflict eruption, and this will be attempted in this subproject.

Feasibility and Risk Management Subproject 5:

The second part of this subproject is very ambitious, but no other subprojects do not depend on this one.

DATA AVAILABILITY AND PROOF OF CONCEPT (PILOT STUDY)

Wikipedia provides all the necessary information in the revision histrory and, more usefully for scraping, in an API which is publicly documented. The API allows to pull the rawtext and metadata on every revision through a unique language specific revision ID. From the rawtext it is also possible to obtain historical information on the presence of langauge links. Articles about the same topic (e.g. battle) are connected through a unique wikidata ID, connecting all language versions. (e.g. https://www.wikidata.org/wiki/Q541626 for the battle of Aspern-Essling). Page views and User profiles are publicly provided by wikimedia's x-tools (https://xtools.wmcloud.org/) which also ensures public data availability through an API.

Since data collection is at the core of this proposal we ran two pilot data collections, with data on 46 and 93 battles in up to 12 languages. We focus on pre-WWI battles to avoid complexities of modern warfare (month-long battles, complex equipment). A focus on human-related casualties ensured a manageable scope, consistent data and kept the dataset aligned with the goals of the pilot. Our larger sample includes over 200,000 Wikipedia articles from 85 languages on 93 historic conflicts spanning the 14th to 20th centuries, edited by 56,000 editors. Among them, 1,721 edited in more than one language, with 76% contributing to English Wikipedia, underscoring its role as an informational aggregation language (Kim et al., 2016). German Wikipedia was second, with 479 multilingual editors (28%).For each conflict, we identified the involved countries and their languages, finding 33 distinct nations, most prominently France. While English Wikipedia covered all 93 conflicts, most other languages had far fewer articles, highlighting the significant barrier language poses to information flow. Future expansions will address underrepresented regions to improve coverage.

To prepare articles for translation and embedding using advanced NLP techniques. Articles are cleaned, fragmented into smaller pieces, and translated into English using Gemini Pro's "text-embedding-004" model, optimized for measuring text relatedness (Team et al., 2023). This process enables clustering, distance, and topic analysis to assess text (dis)similarity. Cosine similarity is applied to derive topic- and language-level averages, identifying the most divergent topics and languages. We also plan to use cluster analysis

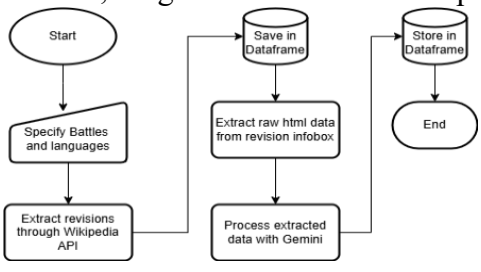


Figure 8: Prototype version of the Page Content

techniques like t-SNE for visualizing variation (Maaten & Hinton, 2008) and apply language fixed effects to account for differences in information density. Future work will attempt using cohere embeddings directly (without pre-translation), providing an alternative approach to ensure robustness of the methodology.

To automate Wikipedia scraping for troop strength and casualty numbers, we use a three-step approach. First, yearly page revisions are collected via the Wikipedia API and saved in a structured data frame. Next, an HTML parser extracts and cleans infobox content for processing. Finally, the Gemini AI processor, combined with Google's translation API, organizes data into structured categories (e.g., troop casualties) with ranges recorded as bounds. AI results are validated against manual checks, with discrepancies resolved through Python postprocessing, ensuring accurate, scalable data collection across languages.

API information, battle boxes and cosine similarity measures were combined in the prototype analysis of battle clusters and narrative distances above. Casualty convergence across 46 battles is also available for analysis.

SCIENTIFIC AND PRACTICAL IMPORTANCE

TruthConverge, ambitious in size and scope, will contribute ground-breaking insights into what drives divergence and reconciliation of diverging versions of 'the truth', in different languages on digital platforms. My strategy will enable me to test widely accepted – but never rigorously tested at scale - notions about how and by whom history is written. It has the potential to disrupt and re-configure basic notions such as the role of winners in historiography and to enrich the methodological toolkit of other disciplines.

In addition to these fundamental contributions, the proposed research offers methodological innovations, by combining newly available research methods based on cross-lingual Large Language Models (AI) with existing rigorous econometric methodology, scraping technologies, and API access to leverage the unique features of Wikipedia as a repository for the construction of language-specific knowledge to generate a novel quantification of the distance between historical narratives across languages. By bridging computational analysis and social science by connecting technical NLP advancements (like LLM-based translation) with social goals and applications (such as conflict resolution, shared historical understanding, we will generate novel understanding of whether (and how) online communities resolve discrepancies in historical accounts of conflict, or even foster reconciliation. This strategy promises to enable boundary-pushing contributions to both fields, Economics and IS as well as History and Memory Studies.

Knowledge Transfer

For greater public benefit, it is necessary to continuously communicate the insights we derive. We will do this through several channels. First, we have already initiated a conversation with the

Wikipedia community and hope to remain in close interaction throughout the project. As a long-term vision, we would hope for the field intervention or the algorithm for detecting increased conflict between narratives to become a standard tool in Wikipedia. Moreover, we will maintain a project website. For outreach with the local community, we will organize two outreach events in the proximity of the workshops draw from the wealth of resources at NovaSBE, where the communication of results to the wider public is a cornerstone of the institution's success.

Risk Analysis and Mitigation Strategies

Several potential risks could affect the successful implementation of this project. These include:

Technical Challenges: Difficulties in obtaining accurate cross-linguistic embeddings and NLP analyses due to intertemporal inconsistencies in Wikipedia content across languages (e.g., varying formats of the "Battle Boxes", that contain information about troop strengths and casualties). Embeddings have already been obtained, and accuracy will be ensured via cross-validation in the data collection using the second measure , research assistants (as requested), and microworking platforms.

Data Limitations: While access to historical edit data is unrestricted, articles for certain languages and conflicts might be subject to incomplete coverage, resulting in missing data and accounts. This risk will be dealt with by including an analysis of content availability as initial part of the analysis. Moreover, we will deploy manual checks for key case studies to validate results from automated analyses.

Intervention Efficacy: The analysis might turn out to show small or limited success of interlinking or AI tools in reducing biases or fostering convergence. For our own field experiment we will use pre-testing and iterative testing methods of the tools in controlled environments before wider implementation will allow to scrutinize the efficiency of the intervention. Structured adjustment to incorporate users' test feedback, combined with additional liaising with the community, and refine the interventions accordingly will ensure maximum efficiency of the intervention.

Stakeholder Engagement: Resistance from Wikipedia editors or platform administrators to adopt proposed tools or methods. To mitigate this risk, we will engage with the Wikipedia community through workshops and forums to obtain their feedback, build trust, and demonstrate the value of the project's tools and findings.

Computational Resource Constraints: High computational costs will be associated with processing large-scale multilingual datasets using advanced AI models. To mitigate the risk of computational resource constraints, funds are sought to ensure sufficient resources and optimize computational processes. Moreover, abilities will further be expanded by leveraging cloud-based AI platforms and distributed computing resources (for aggregate or anonymized data to ensure the privacy protection of any involved individuals).

If certain risks materialize nonetheless, contingency plans are in place, which include:

For data limitations: Narrow the project scope to fewer, more thoroughly analyzed case studies without compromising the research goals. **For technical challenges** we have ensured basic data availability and can adjust research focus to highlight qualitative insights from the available data rather than quantitative metrics alone. Generally, each subproject is designed with increasing levels of risk, so that every subproject's scope can be adjusted if needed. **For stakeholder engagement**, we will publish open-access tools and methods to allow gradual adoption at the community's discretion.

Expected Impact and Outlook.

By developing innovative measurement strategies, the project seeks to identify areas of narrative divergence, highlighting the **potential of online platforms** like Wikipedia **to bridge gaps** in understanding **- rather than divide** - and foster a process of **converging 'truths**.' Additionally, it will examine the factors that moderate this process and assess the impact of targeted interventions, such as flagging diverging narratives, to promote alignment. On the other hand, non-convergence or even growing divergence in narratives could turn out to be a valuable indicator of latent conflict or even enable early detection of future outbreaks.

The **broader impact** of this work lies in its **emphasis on fostering a mutual and shared understanding of historical events**, which is crucial for conflict resolution and prevention. By quantifying the "heat" of ongoing conflicts and potentially creating an **early warning system for impending disputes**, the research aspires to make a **high-stakes contribution** to global peacebuilding efforts.

8. References

Adafre, S. F., & de Rijke, M. (2006). Finding similar sentences across multiple languages in Wikipedia. *Proceedings of the EACL 2006 Workshop on New Text*, 62–69.

Adar, E., Skinner, M., & Weld, D. S. (2009). Information arbitrage across multilingual Wikipedia. *Proceedings of the Second ACM International Conference on Web Search and Data Mining*, 94–103.

Álvarez, G., Oeberst, A., Cress, U., & Ferrari, L. (2020). Linguistic evidence of in-group bias in English and Spanish Wikipedia articles about international conflicts. *Discourse, Context & Media, 35*, 100391. https://doi.org/10.1016/j.dcm.2020.100391

Araujo, A. L. (Ed.). (2013). *Politics of memory: Making slavery visible in the public space.* Routledge.

Assmann, A. (2011). *Cultural memory and Western civilization: Functions, media, archives.* Cambridge University Press.

Assmann, A., & Conrad, S. (Eds.). (2010). *Memory in a global age: Discourses, practices, and trajectories.* Palgrave.

Barberá, P. (2020). Social media, echo chambers, and political polarization. In N. Persily & J. A. Tucker (Eds.), *Social media and democracy: The state of the field, prospects for reform* (p. 34). Cambridge University Press.

Bekerman, Z., & Zembylas, M. (2011). *Teaching contested narratives: Identity and memory in peace education.* Cambridge University Press.

Blanshard, B. (1939). *The Nature of Thought* (Vols. 1–2). George Allen & Unwin.

Bloch, M. (1954). *The historian's craft.* Vintage Books.

Buckley-Zistel, S. (2006). Remembering to forget: Chosen amnesia as a strategy for local coexistence in post-genocide Rwanda. *Africa: Journal of the International African Institute, 76*(2), 131–150.

Carr, E. H. (1961). *What is history?* Macmillan.

Callahan, E. S., & Herring, S. C. (2011). Cultural bias in Wikipedia content on famous persons. *Journal of the American Society for Information Science and Technology, 62*(10), 1899–1915.

Callaway, B., & Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics, 225*(2), 200–230.

Choi, D., Chun, S., Oh, H., Han, J., & Kwon, T. T. (2020). Rumor propagation is amplified by echo chambers in social media. *Scientific Reports, 10*(1), 310.

Cinelli, M., De Francisci Morales, G., Galeazzi, A., Quattrociocchi, W., & Starnini, M. (2021). The echo chamber effect on social media. *Proceedings of the National Academy of Sciences, 118*(9).

Clark, P. (2010). *The Gacaca courts, post-genocide justice and reconciliation in Rwanda: Justice without lawyers.* Cambridge University Press.

Cohere. (n.d.). *Cohere embeddings*. Retrieved 10/12/2024, from https://cohere.com

Cornelißen, C. (2006). Europas Gedächtnislandkarte: Gibt es eine Universalisierung des Erinnerns? In N. Frei (Ed.), *Was heißt und zu welchem Ende studiert man Geschichte des 20. Jahrhunderts?* (pp. 42–49). Wallstein.

Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2018). BERT: Pre-training of deep bidirectional transformers for language understanding. *CoRR*, abs/1810.0.

Doxtader, E. (2009). *With faith in the works of words: The beginnings of reconciliation in South Africa.* Michigan State University Press.

Erdmann, M., Nakayama, K., Hara, T., & Nishio, S. (2009). An approach for extracting bilingual terminology from Wikipedia. *Proceedings of the 13th International Conference on Database Systems for Advanced Applications*, 380–392.

Erll, A. (2011). *Memory in culture.* Palgrave Macmillan.

Fallis, D. (2008). Toward an epistemology of Wikipedia. *Journal of the American Society for Information Science and Technology, 59*(10), 1662–1674.

Ford, H., Sen, S., Musicant, D. R., & Miller, N. (2013). Getting to the source: Where does Wikipedia get its information from? *Proceedings of the 9th International Symposium on Open Collaboration*, 9:1–9:10.

Forte, A., & Bruckman, A. (2008). Scaling consensus: Increasing decentralization in Wikipedia governance. *Proceedings of the 41st Annual Hawaii International Conference on System Sciences*, 157.

Forte, A., Lampe, C., & Bruckman, A. (2012). Scaling consensus: Collective editing in Wikipedia. *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work*, 161–170.

Forte, A., & Lampe, C. (2013). Defining, understanding, and supporting open collaboration: Lessons from the literature. *American Behavioral Scientist, 57*(5), 535–547.

Galtung, J. (1996). *Peace by peaceful means: Peace and conflict, development and civilization.* SAGE Publications.

Garimella, K., De Francisci Morales, G., Gionis, A., & Mathioudakis, M. (2018). Political discourse on social media: Echo chambers, gatekeepers, and the price of bipartisanship. In *Proceedings of the 2018 World Wide Web Conference* (pp. 913–922).

Garde-Hansen, J. (2011). *Media and memory*. Edinburgh University Press.

Garton Ash, T. (2011, March 16). Germany can show reborn Arab nations the art of overcoming a difficult past. *The Guardian.*

Greenstein, S., & Zhu, F. (2012). Is Wikipedia biased? *American Economic Review, 102*(3), 343–348. https://doi.org/10.1257/aer.102.3.343

Greenstein, S., & Zhu, F. (2018). Do experts or crowd-based models produce more bias? Evidence from Encyclopædia Britannica and Wikipedia. *MIS Quarterly, 42*(3), 945–959. https://doi.org/10.25300/MISQ/2018/11432

Habermas, J. (1984). *The Theory of Communicative Action: Reason and the Rationalization of Society* (Vol. 1). Beacon Press.

Halatchliyski, I., Kummer, M. E., Slivko, O., Giorgidze, G., & Saam, M. (2016). Centrality and content creation in networks: The case of economic topics on German Wikipedia. *Information Economics and Policy, 36*, 36–52. https://doi.org/10.1016/j.infoecopol.2016.05.001

Halbwachs, M. (1992). *On collective memory*. University of Chicago Press.

Hale, S. A. (2014). Multilinguals and Wikipedia editing. *Proceedings of the 2014 ACM Conference on Web Science*, 99–108.

Hayner, P. B. (2010). *Unspeakable truths: Transitional justice and the challenge of truth commissions.* Routledge.

Hecht, B., & Gergle, D. (2010). The Tower of Babel meets Web 2.0: User-generated content and its applications in a multilingual context. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 291–300.

Hinnosaar, M., Hinnosaar, T., Kummer, M., & Slivko, O. (2023). Wikipedia matters. *Journal of Economics & Management Strategy*. https://doi.org/10.1111/jems.12421

Hunt, L. (1989). *The new cultural history*. University of California Press.

James, W. (1907). *Pragmatism: A New Name for Some Old Ways of Thinking.* Longmans, Green, and Co.

Jemielniak, D. (2014). *Common Knowledge? An Ethnography of Wikipedia.* Stanford University Press.

Jenkins, H. (2006). *Convergence Culture: Where Old and New Media Collide.* New York University Press.

Kim, S., Park, S., Hale, S. A., Kim, S., Byun, J., & Oh, A. (2016). Understanding editing behaviors in multilingual Wikipedia. *PLOS ONE, 11*(5). https://doi.org/10.1371/journal.pone.0155305

Kirkham, R. L. (2001). *Theories of Truth – A Critical Introduction.* MIT Press.

Knigge, V. (2008). Europäische Erinnerungskultur: Identitätspolitik oder kritisch-kommunikative historische Selbstvergewisserung. In Kulturpolitische Gesellschaft e.V. (Ed.), *kultur.macht.europa – europa.macht.kultur* (pp. 69–82). Klartext.

Kittur, A., & Kraut, R. E. (2008). Harnessing the wisdom of crowds in Wikipedia: Quality through coordination. *Proceedings of the 2008 ACM Conference on Computer Supported Cooperative Work*, 37–46.

Kummer, M., Slivko, O., & Zhang, X. (2019). Unemployment and digital public goods contribution. *Information Systems Research.* https://doi.org/10.1287/isre.2019.0916

Laitenberger, U., Viete, S., Slivko, O., Kummer, M., Borchert, K., & Hirth, M. (2023). Unemployment and online labor: Evidence from microtasking. *MIS Quarterly.* https://doi.org/10.25300/MISQ/2022/17347

Larson, D. W., & Shevchenko, A. (2010). Status seekers: Chinese and Russian responses to US primacy. *International Security, 34*(4), 63–95.

Lederach, J. P. (2005). *The moral imagination: The art and soul of building peace.* Oxford University Press.

Leggewie, C., & Lang, A. (2011). *Der Kampf um die europäische Erinnerung: Ein Schlachtfeld wird besichtigt.* C.H. Beck.

Lewoniewski, W., Węcel, K., & Abramowicz, W. (2017). Multilingual ranking of Wikipedia articles with quality and popularity assessment in different language editions. *Online Information Review, 41*(4), 478–495.

Lih, A. (2009). *The Wikipedia Revolution: How a Bunch of Nobodies Created the World's Greatest Encyclopedia.* Hyperion.

Liu, S., Yan, X., Li, H., Zhao, T., & Wang, Y. (2018). Cross-lingual entity alignment via attribute-enhanced embedding. *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 2486–2495.

Luyt, B., & Tan, D. (2010). Improving Wikipedia's accuracy: Is edit age a solution? *Journal of the American Society for Information Science and Technology, 61*(12), 2582–2593.

Maaten, L. van der, & Hinton, G. (2008). Visualizing data using t-SNE. Journal of Machine Learning Research, 9(86), 2579–2605.
Magnus, P. D. (2009). On trusting Wikipedia. *Episteme, 6*(1), 74–90.
Minow, M. (1998). Between vengeance and forgiveness: Facing history after genocide and mass violence. Beacon Press.
Miller, D. L. (2023). An introductory guide to event study models. *Journal of Economic Perspectives,* 37(2), 203-230.
Miquel-Ribé, M., & Laniado, D. (2018). Wikipedia Cultural Diversity Dataset: A cartography of Wikipedia and Wikimedia projects. *Proceedings of the 14th International Symposium on Open Collaboration*, 1–9.
Montville, J. V. (1990). The healing function in political conflict resolution. Journal of Social Issues, 46(2), 33–53.
Noddings, N., & Brooks, L. (2016). Teaching controversial issues: The case for critical thinking and moral commitment in the social studies. Teachers College Press.
Nora, P. (1996). Realms of memory: Rethinking the French past. Columbia University Press.
Olick, J. K., Vinitzky-Seroussi, V., & Levy, D. (2011). The collective memory reader. Oxford University Press.
Peirce, C. S. (1878). How to make our ideas clear. *Popular Science Monthly, 12*, 286–302.
Pfeil, U., Zaphiris, P., & Ang, C. S. (2006). Cultural differences in collaborative authoring of Wikipedia. *Journal of Computer-Mediated Communication, 12*(1), 88–113.
Radonic, L. (2010). Krieg um die Erinnerung: Kroatische Vergangenheitspolitik zwischen Revisionismus und europäischen Standards. Campus.
Radonic, L. (2014). Postsozialistische Gedenkmuseen zwischen Opfernarrativen und der 'Europäisierung der Erinnerung.' Jahrbuch für Politik und Geschichte, 5, 85–106.
Rambachan, A., & Roth, J. (2019). An honest approach to parallel trends. *Working paper, Harvard University*.
Rambachan, A., & Roth, J. (2023). A more credible approach to parallel trends. *Review of Economic Studies*, *90*(5), 2555-2591.
Reagle, J. (2010). *Good Faith Collaboration: The Culture of Wikipedia.* MIT Press.
Regan, P. (2010). *Unsettling the settler within: Indian residential schools, truth telling, and reconciliation in Canada.* UBC Press.
Rescher, N. (1973). *The Coherence Theory of Truth.*
Rescher, N. (1993). *A System of Pragmatic Idealism, Volume II: The Validity of Values.* Princeton University Press.
Ricoeur, P. (2004). *Memory, history, forgetting.* University of Chicago Press.
Rousso, H. (2004). Das Dilemma eines europäischen Gedächtnisses. *Zeithistorische Forschungen*, *1*, 363–378.
Rousso, H. (2011). History of memory, policies of the past: What for? In K. H. Jarausch & T. Lindenberger (Eds.), *Conflicted memories: Europeanizing contemporary histories* (pp. 23–38). Berghahn.
Samoilenko, A., Karimi, F., Edler, D., Kunze, T., & Strohmaier, M. (2017). Linguistic neighbourhoods: Explaining cultural borders on Wikipedia through language co-editing activity. *Proceedings of the Eleventh International AAAI Conference on Web and Social Media*, 217–226.
Simpson, A. (2014). *Mohawk interruptus: Political life across the borders of settler states.* Duke University Press.
Sommer-Sieghart, M. (2006). Historische Ausstellungen als 'Contested Space.' In J. Feichtinger et al. (Eds.), *Schauplatz Kultur – Zentraleuropa: Transdisziplinäre Annäherungen* (pp. 159–166). Studienverlag.
Team, Gemini, Anil, R., Borgeaud, S., Wu, Y., Alayrac, J.-B., Yu, J., Soricut, R., Schalkwyk, J., Dai, A. M., Hauth, A., Millican, K., Silver, D., Petrov, S., Johnson, M., Antonoglou, I., Schrittwieser, J., Glaese, A., Chen, J., Pitler, E., ... & Vinyals, O. (2023). Gemini: A family of highly capable multimodal models. arXiv preprint arXiv:2312.11805.
Terren, L. T. L., & Borge-Bravo, R. B. B. R. (2021). Echo chambers on social media: A systematic review of the literature. *Review of Communication Research, 9*.
Tosh, J. (2015). *The pursuit of history.* Routledge.
Tutu, D. (1999). *No future without forgiveness.* Random House.
Vrandecic, D. (2012). Wikidata: A new platform for collaborative data collection. *Proceedings of the 21st International Conference on World Wide Web*, 1063–1064.
Weizman, E. (Ed.). (2014). *Forensis: The architecture of public truth.* Sternberg Press.
Wilson, R. (2001). *The politics of truth and reconciliation in South Africa: Legitimizing the post-apartheid state.* Cambridge University Press.
Wooldridge, J. M. (2010). *Econometric Analysis of Cross Section and Panel Data* (2nd ed.). MIT Press.
Zehr, H. (2002). *The little book of restorative justice.* Good Books.

Zhang, X. (Michael), & Zhu, F. (2011). Group size and incentives to contribute: A natural experiment at Chinese Wikipedia. *American Economic Review, 101*(4), 1601–1615.

Zhu, K., Walker, D., & Muchnik, L. (2020). Content growth and attention contagion in information networks: Addressing information poverty on Wikipedia. *Information Systems Research, 31*(2), 491–509.

To measure how narratives change and evolve over time, we will leverage the unique features of Wikipedia that provide us with a repository of knowledge items (such as the page about a past battle) that are cross-referenced across up to 80 different languages. Moreover, as the entire edit history has been recorded and is publicly available, we can study the evolution of each article in each language over time and at any desired frequency (monthly, hourly), to measure variables that characterize editing processes (information availability, editing density, contentedness) and article qualities (length, references, readability, etc.).

We envision a feedback loop in which the initial analysis is to inform the fine-tuning and calibration of the quantification methodology. Hence, we will focus on historic battles that are well documented in several languages and were fought over 175 years ago. Not only are these battles ideally suited for the research questions, but they also provide us with a second metric based on cross-language divergence in key numbers and statistics regarding troop strength and casualties on both sides. Pages about uncontested content will serve as comparison group. Several suitable page-types are available, including content about geometric objects or celestial bodies, that allow to control for baseline differences between the languages under study.