


RESEARCH ARTICLE

Algorithmic influence on conflict and cooperation in digital communities

Andrea Russo ^{1,2*}, Alessandro Caliandro³

1 University of Pavia, Sociology Department, Pavia, Italy, **2** University of Catania, Department of Physics, Catania, Italy, **3** University of Milan, Department of Social and Political Sciences, Milan, Italy

* Andrea.russophd@gmail.com



Abstract

This study focus on the contrasting dynamics of discussion and information dissemination on two influential platforms: X and Wikipedia. While X often serves as a battleground for contentious debates, where users engage in direct confrontation, Wikipedia fosters a collaborative environment aimed at reaching consensus. Focusing on polarizing issues such as the ongoing Russian invasion of Ukraine, the research examines how information is shared, contested, and shaped within these distinct communities. Data from X, collected using the hashtag #UkraineRussiaWar, undergoes a NLP process to categorize and highlight the diverse topics discussed. By classifying the data into Pro-war or Against-war perspectives, I analyze user reactions and the primary themes driving these discussions. For Wikipedia, I gathered and analyzed comments and contributions from various authors, providing insights into the collaborative discourse and consensus-building process.

OPEN ACCESS

Citation: Russo A, Caliandro A (2026) Algorithmic influence on conflict and cooperation in digital communities. PLOS Complex Syst 3(4): e0000087. <https://doi.org/10.1371/journal.pcsy.0000087>

Editor: Hocine Cherifi, Université de Bourgogne: Universite de Bourgogne, FRANCE

Received: April 14, 2025

Accepted: December 20, 2025

Published: April 6, 2026

Peer Review History: PLOS recognizes the benefits of transparency in the peer review process; therefore, we enable the publication of all of the content of peer review and author responses alongside final, published articles. The editorial history of this article is available here: <https://doi.org/10.1371/journal.pcsy.0000087>

Copyright: © 2026 Russo, Caliandro. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and

Author summary

This study investigates how two major digital platforms, X and Wikipedia—shape discussions and manage conflict regarding the ongoing Russia-Ukraine war. While X often fosters polarized debates, Wikipedia promotes collaborative editing aimed at consensus. By analyzing user-generated content, I uncover distinct communication patterns and the role of platform design in shaping discourse. On X, communities were divided into pro-war and anti-war factions, each focusing on different narratives, with minimal cross-group dialogue. The pro-war group emphasized geopolitical arguments, while the anti-war side highlighted humanitarian concerns. This lack of a shared topic further undermined collaboration, as discussions became fragmented, and by consequence reducing opportunities for mutual understanding. In contrast, Wikipedia discussions centered on factual accuracy, naming conventions, and evidence-based edits, reflecting a more structured and less polarized environment. Using natural language processing (NLP) techniques, it was found that X's discourse exhibited lower entropy, indicating simpler

reproduction in any medium, provided the original author and source are credited.

Data availability statement: The data for this manuscript can be found at <https://zenodo.org/records/14937452>.

Funding: This work was financially supported by: European Union - Next Generation EU, Missione 4, Componente 1, CUP: G53D23004770006, Project name: Feedback culture: assessing the effects of algorithmic recommendations on platformized consumption (ALGOFEED).

Competing interests: The authors have declared that no competing interests exist.

language, while Wikipedia's higher entropy reflected more complex, nuanced exchanges. Our findings highlight how platform structures and community norms influence the dynamics of online conflict and cooperation, offering insights for designing healthier digital ecosystems that foster productive dialogue and knowledge sharing.

1 Introduction

The digital age has profoundly transformed the way information is generated, shared, and assimilated. Collaborative platforms like Wikipedia, alongside social media such as X (previously known as X), have become central to the global information ecosystem, driving both knowledge/information production and online social dynamics. Understanding the interactions, conflicts, and collaborative dynamics on these platforms not only sheds light on important aspects of digital culture but also provides valuable tools for understanding how platform design influence the way people discuss and resolve conflicts.

In this context, the analysis of user interactions offers a deeper insight into how collective behavior can shape emerging and broader cultural phenomena.

Wikipedia and X are platforms that have already been compared for cooperation and content moderation cases [1], while other have focused to explored the intricacies interaction and reality of these platforms like the use of Wikipedia users activity as a early predicting method for box office success [2], highlighting the predictive power of big data in cultural contexts. But also several study [3] have focused on the dynamics of conflicts within Wikipedia, revealing the complex nature of editorial disputes and their resolution processes within the platform. This work is further previously supported by studies on "edit wars" [4], which delve into the patterns of conflict and cooperation among Wikipedia editors.

Moreover, [5] provided a cross-disciplinary survey on human-machine networks, offering insights into the broader implications of these interactions across various platforms, including Wikipedia and social media. In a similar vein, research on echo chambers and filter bubbles by [6] and [7] explored how these phenomena can significantly impact the dissemination and perception of information on social media.

The role of Wikipedia in the broader context of online collaboration and conflict has been examined from different angles, evidencing how collaborative efforts contribute to the platform's overall reliability [8] or, from the opposite side, highlighting the factors that cause disputes in online knowledge communities [9], providing an empirical basis for understanding the dynamics of topic conflicts. Visualizations of these conflicts have been studied by [10], who used history flow visualizations to track cooperation and conflict among Wikipedia authors. Furthermore, the relationship between conflict and language complexity has been explored in the context of Wikipedia by [11], who found that controversial articles tend to reduce language complexity on their associated talk pages. This finding suggests that conflict may not only affect the collaborative dynamics but also the linguistic features of discourse within the platform.

In addition to Wikipedia, studies have examined the interplay between

Wikipedia and other platforms such as the model [12] that thanks to the combination of Google, Wikipedia, and X data could lead as an indicator of coronavirus deaths, highlighting the interconnectedness of these platforms [13] in tracking real-world events, also in localized phenomena [14]. Lexical comparisons between Wikipedia and X have also been studied, with [15] using word embeddings to compare the two corpora, shedding light on the linguistic differences and similarities across platforms, while [16] have used social networks within Wikipedia itself to examine the structure and dynamics of social interactions among Wikipedia contributors. Always through the use of NLP tools, [17–19] investigated the role of emotions in the spread of misinformation on social media, finding that emotionally charged content is more likely to be shared, which can exacerbate the spread of misinformation and an influence from the social-cultural environment that in the end significantly influence collective behavior.

Therefore, for platforms with a high rate of polarisation and ‘conflict’ (understood as a difference of opinion on a specific subject) such as X, studies have further deepened our understanding of these dynamics by focusing on the role of algorithms and misinformation, in the political sphere [20] highlighting how algorithms can influence the spread of political information, potentially exacerbating polarization, later reinforced by research [21] that conducted an audit of algorithmic bias on X, uncovering patterns that suggest systemic biases in how content is presented to users, complemented also from other work [22], that investigated the impact of fake news on X during the 2016 U.S. presidential election, revealing how misinformation can spread rapidly and influence public opinion. Moreover, polarisation and conflict are not only present on X, in fact it was examined how these cases are also present on ‘Collaborative’ platforms [23], by examining the emergence and regulation of extremist behavior.

These studies collectively underscore the importance of analyzing both collaborative and conflictual dynamics on platforms like Wikipedia and X, as well as the algorithmic influences that shape the information landscape. Together, they offer a window into the broader digital environment, where information dissemination, public opinion, and social behavior are increasingly interlinked, and where algorithms play a pivotal role in shaping the flow of information.

While much research has focused on algorithms, dynamics, and the environments where people discuss and share information, there is a notable lack of studies analyzing the resolution and management of conflicts that arise during contentious discussions. This work aims to address this gap.

2 Results

To explore the conflict and cooperation dynamics in different platforms, I first ask: How can we detect and measure cooperation dynamics in a given network?

I came up with a textual analysis to link the interactions between the various accounts, and to measure the principles of conflict and/or cooperation through the words used. To get the most relevant words within each cluster, I extracted the most relevant topics for each interaction using Bigram, so as to also link people to the discussion topics.

The macro discussion topic is Russia’s invasion of Ukraine in 2022.

Figs 1 and 2 represent the two different X communities, the Pro-War and Against-War obtained. The different colours represent the communities of words on a given topic. It can already be seen that the Pro-war community, i.e. the people who share the idea that the invasion of Ukraine is correct according to their reasons shows how certain topics such as presented in Table 1.

While Against war express their contrariety against the war using other kind of arguments, show in Table 2.

One notices how the pro-war group used the excuse of the NATO presence and the hypothetical 2014 ‘Revolution of Dignity’ coup d’état as main topics, while the anti-war community marginalises the NATO topic by giving more attention to the suffering of people under attack and defining Putin as a war criminal.




Fig 1. Pro War topics network.

<https://doi.org/10.1371/journal.pcsy.0000087.g001>

In order to establish a relationship of resolution and cooperation, it is necessary to create at least a single and shared channel for both groups, and the only one present as an exploitable topic to start a discourse is that of NATO, which is present in both communities, but there is a strong evidence of polarisation on the subject.

Fig 3 shows the topics and discussions within the Wikipedia community. Unlike the X community, Wikipedia has the need to find a common solution for all accounts (both pro-war and anti-war), and in fact it is noticeable that none of the previous X topics were used. Most of the discussions and interactions between users are focused on different topics as show in Table 3.

Most of the discussions are obviously editing requests, reading and interactions between the various accounts on the choices to be made on certain topics (light blue) and discussions between the various accounts (dark teal).

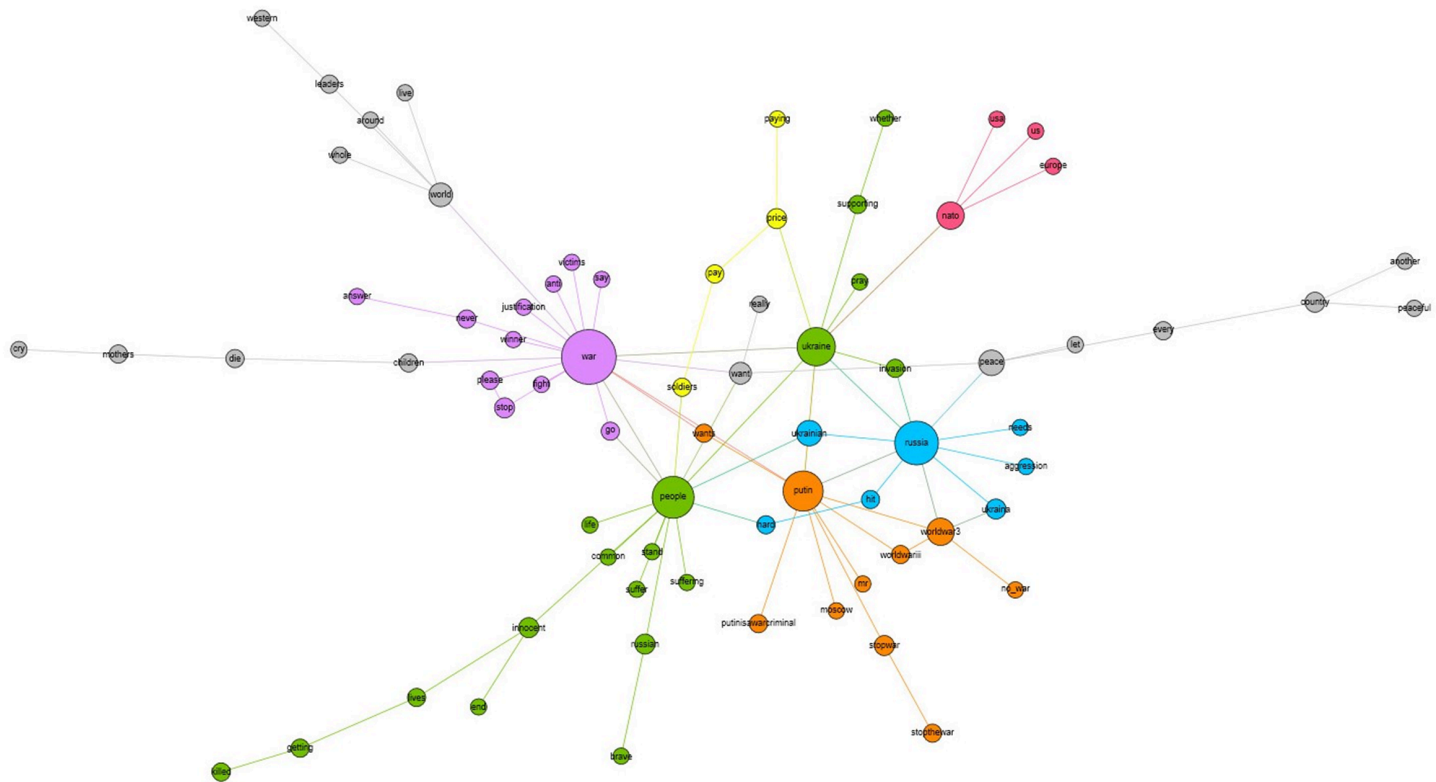


Fig 2. Against War topics network.

<https://doi.org/10.1371/journal.pcsy.0000087.g002>

Table 1. Topics and corresponding community colours (Pro-War).

Topics	Color Community
2014 coup d'état that installed a puppet regime	Bordeaux
NATO expansion and presence close to Russia	Orange
US soldiers' crimes	Green
European nations should ally with the 'Soviets' rather than European governments	Grey

<https://doi.org/10.1371/journal.pcsy.0000087.t001>

Table 2. Topics and corresponding community colours (Against-War).

Topics	Color Community
Innocent people are being killed	Green
Putin is an international criminal	Orange
Children die and mothers cry	Grey
Soldiers pay the price (with death)	Yellow

<https://doi.org/10.1371/journal.pcsy.0000087.t002>

Starting from these initial results, it is possible to see how the two communities are different according to the topics discussed. For example, Wikipedia discussions did not particularly focus on the subject of NATO (I assume since there are no references or evidence), while X do. On the other hand, however, the absence of sources on the part of the pro-war community and the lack of sources on the part of the anti-war community highlights how the content of the discussions does not present a collaborative environment, as everyone expresses his or her opinion without relying on sources,

Table 4. Summary of perspectives and categorization.

Category	Examples	Type
Pro-War	2014 coup d'état that installed a puppet regime NATO expansion and presence close to Russia US soldiers' crimes European nations should ally with the 'Soviets' rather than European governments	Ideological and geopolitical
Against-War	Innocent people are being killed Putin is an international criminal Children die and mothers cry Soldiers pay the price (with death)	Emotional and humanitarian
Wikipedia	The list of countries supporting Ukraine Russia uses armed separatist forces Changing the event name: ' <i>Russian-Ukrainian-War</i> ' to ' <i>Russian invasion of Ukraine</i> '	Technical/ Organisational and descriptive

<https://doi.org/10.1371/journal.pcsy.0000087.t004>

Meanwhile, the *Wikipedia community* adopts a distinctly technical, organizational, and descriptive approach. Contributions here center on factual clarifications, such as updating conflict terminology, identifying the actors involved, and ensuring contextual accuracy.

While X fosters ideological and emotional narratives, knowledge-driven platforms like Wikipedia prioritize objective documentation and descriptive clarity. Regarding the cooperative dynamics, this table evidence how it is already difficult, especially for the two X groups, given the absence of a prevailing topic in common, to try to interact with the aim of discussing and confront each other in order to collaborate and find common ground.

This fragmentation not only reflects the different priorities of each community but could also highlight the broader challenge of fostering dialogue across platforms characterized by contrasting communicative approaches. The lack of thematic overlap could reduce opportunities for meaningful exchange, further entrenching communities within their respective discursive boundaries and potentially reinforcing echo chambers rather than promoting cross-community engagement.

Fig 4 represents the activity of the various accounts that have collaborated to edit, update and improve the Wikipedia page.

This heatmap reveals two notable observations. First, it highlights that activity was already present well before the onset of the Russian invasion of Ukraine, specifically from two accounts: *Tobby72* and *8.28.81.21*. Both accounts had previously engaged with the Wikipedia community, requesting that the page title be modified to “Russo-Ukrainian War” before to the actual invasion by Russia; This could be part of the Russian propaganda that had already decided to invade Ukraine many months early. The second observations is that there are only a few accounts that expound their ideas once and are characterised by the fact that they do not have an editor’s account name (the names are expressed in numbers given by Wikipedia itself to distinguish them from each other), while the ‘professional’ accounts are the most active ones who usually have to correct and prompt other accounts for modification on the page or suggested changes. However, a potential limitation is that a share of X and Wikipedia accounts might, in principle, belong to the same individual. Despite the relatively limited size of the datasets considered, it is not possible, either through the X API or through Wikipedia scraping, to obtain account-level identifiers such as email addresses or other information that would allow for verification and cross-platform matching.

Moreover, even if such information were hypothetically available, the same individual could still use different accounts (including distinct email addresses) on different platforms. As a result, any matching procedure would carry a high risk of false negatives, as there is a possibility that the same individual could use two different email addresses.

Fig 5 illustrates the network interactions among Wikipedia accounts, providing information by incorporating both node and edge attributes.




Fig 4. Wikipedia accounts activity.

<https://doi.org/10.1371/journal.pcsy.0000087.g004>

The nodes are represented by their relative importance (size) and community affiliation (colour), while the edges convey additional layers of information, including the topics discussed (extracted through BERTopic), sentiment, emotion, and Shannon’s entropy. This graph reveals that certain accounts have played a significantly more prominent role than others, demonstrating higher levels of activity in editing, commenting, and proposing changes to the Wikipedia page. Moreover, it underscores how interactions among users can differ not only in terms of entropy but also in the emotional tone they convey. Lastly, many grey-coloured accounts are characterized solely by self-loops, indicating that they proposed changes to the text without receiving any response. In such cases, I assume that the lack of response implies acceptance of the suggested modifications.

A dynamic, animated version can be found at the above [Link](#).

However a substantial difference obtained by analysing the differences between Sentiment and Entropy can be seen in [Figs 6 and 7](#).

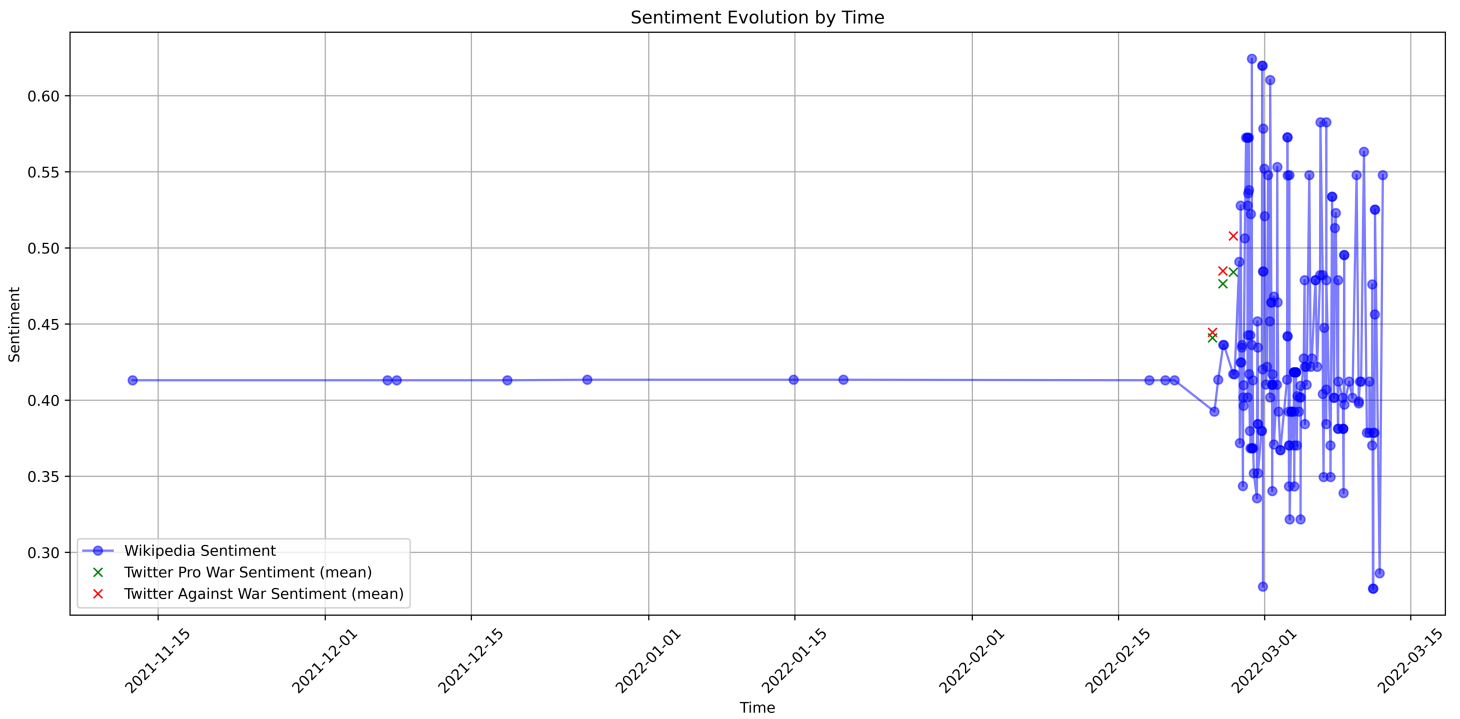


Fig 6. X/Twitter and Wikipedia sentiment evolution by time.

<https://doi.org/10.1371/journal.pcsy.0000087.g006>

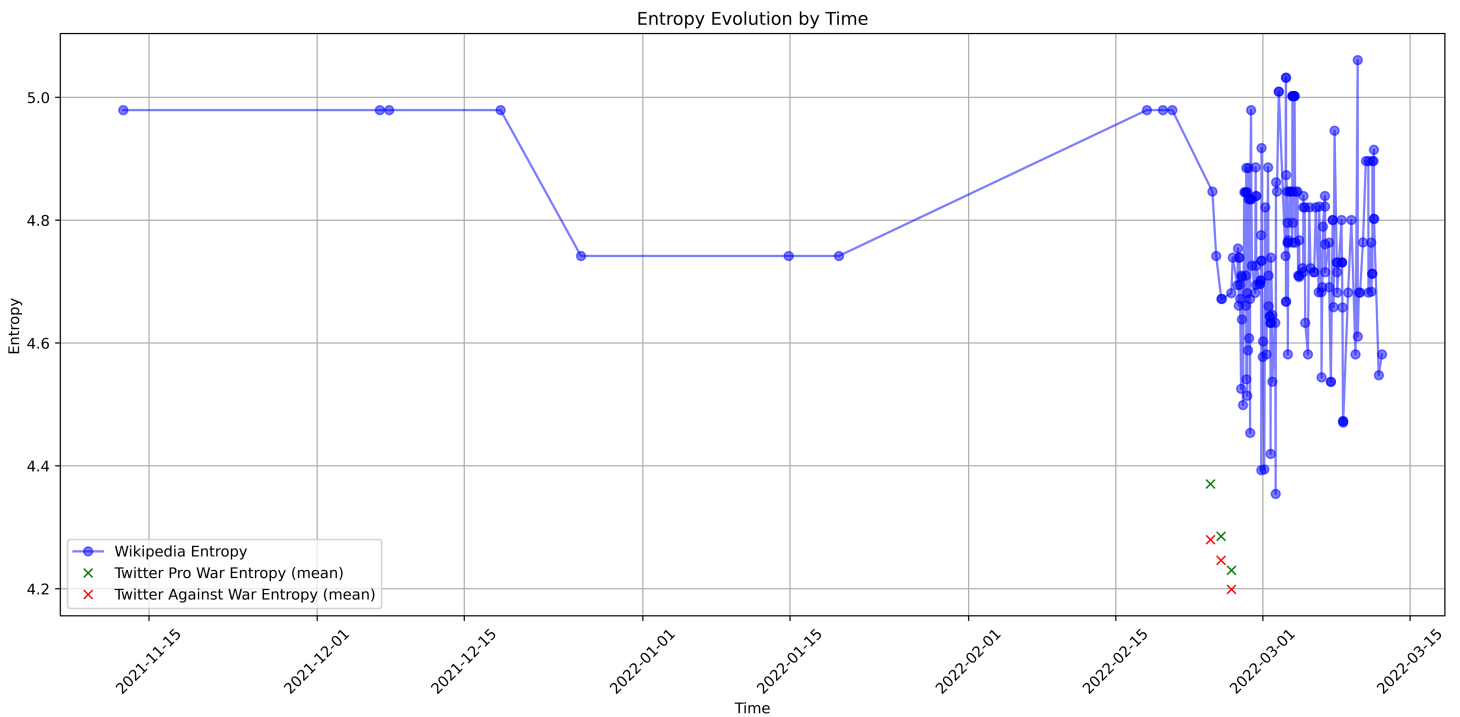


Fig 7. X/Twitter and Wikipedia entropy evolution by time.

<https://doi.org/10.1371/journal.pcsy.0000087.g007>

Wikipedia. This difference can be attributed to the nature of content creation and the communicative purposes that each platform serves.

On Wikipedia, contributions are expected to be supported by credible sources, and editors are often required to use precise, formal, and contextually rich language to justify their edits and ensure adherence to the platform's standards of verifiability and neutrality. The necessity for such a level of rigour and detail could result in higher entropy, reflecting a more complex linguistic structure. A critical aspect of the Wikipedia community is its unique platform culture, which creates the illusion of both 'work' and 'volunteering.' While contributors engage voluntarily, the structured nature of the platform, including editorial guidelines, peer review processes, and dispute resolution mechanisms, mirrors organized labor, creating a collaborative effort by the social environment, which helps the proliferation and growth of collaboration.

While, X's character limit and its role as a platform for rapid information sharing and real-time reactions naturally favour more straightforward language, minimizing complexity to maximize engagement and immediacy, Wikipedia's structured environment encourages dialogue, compromise, and the pursuit of neutrality, facilitating more constructive forms of cooperation.

These contrasting cultural approaches also shape the dynamics of conflict and cooperation within each platform. One key tension lies in the social rules and objectives that govern user participation. On X, users act as "citizens," exercising free speech with minimal constraints, whereas on Wikipedia, they function as "volunteers," adhering to community guidelines and collaborative standards.

3 Discussion

The comparative analysis of X and Wikipedia during the ongoing conflict between Russia and Ukraine offers significant understanding into how information is shared, discussed, and shaped on different platforms. By examining the Pro- and Against-war networks on X and the discussions within the Wikipedia community, I identified stark differences in how topics are addressed, the complexity of language used, and the overall interaction dynamics among users. While Emotion detection and Sentiment analysis did not yield significant results, the Topic and Entropy did, but a great importance comes from the approach to the platform. On X, conversations are marked by polarized communities, each focused on specific narratives—such as NATO's involvement or humanitarian concerns—without much overlap. The pro-war community tends to use historical and geopolitical arguments, often devoid of robust sourcing, while the anti-war community focused around human rights and condemnation of war crimes. This polarization, combined with the platform's brevity and immediacy, leads to a lower entropy in the language used, favouring simple, direct statements that can quickly resonate or provoke reactions. Conversely, Wikipedia discussions illustrate a more structured and moderated environment, where the necessity of achieving consensus and maintaining neutrality is paramount. The topics discussed are less about taking a side and more about accurately framing events, naming conventions, and presenting verified information. The requirement for citations and evidence fosters a higher entropy in the language, as contributors must craft more precise and complex arguments to validate their edits. This is further evidenced by the observed collaborative patterns, where only a few dedicated editors frequently engage in discourse to maintain the page's integrity, while casual or one-time users have a limited impact. Another critical aspect to consider is how contributors on Wikipedia often perceive their involvement as either work or volunteering. This cultural perspective could tend to encourage a more constructive and articulated approach over time, which could significantly influence the dynamics of conflict and cooperative behaviour within the community.

4 Materials and methods

All the tests and results of this work can be found at the [Link](#), including a temporal animation of Wikipedia user interactions.

In this work, I analyse the interactions between X and Wikipedia platforms to observe how and if, solutions and cooperation can be found between users in the respective social networks. As a ‘battlefield’, the topic of the Russia-Ukrainian war/invasion were chose, which is still in the news today two years after (2022/02/24) the invasion.

While Wikipedia remains a free platform for data collection, X has undergone different changes in recent years, especially since its new acquisition. The data extracted and present in this publication were taken prior to new X’s policy change.

Data were collected from both X (with the hashtag #RussiaUkraineConflict and #RussiaUkraineWar) (Chosen as the main hashtags with the highest number of interactions in those days) and Wikipedia during the invasion of Ukraine by Russian forces in 2022/02/24, but two days after the invasion, X was blocked by the Russian government. Due to the inability and limitation of the Russian population to express opinions on the ongoing invasion I deemed it necessary to use only the data from the first two days, resulting in a total of 5,480 tweets. As for Wikipedia data, users interactions from one month before the war and one month after the war were downloaded, for a total of ~ 85 discussion topics.

Collaborative networks were created for both databases, but while for Wikipedia it is common to interact with several people for a given discussion and thus giving the possibility to create a network, for X it was challenging to identify a cohesive network of user interactions, as many participants preferred to communicate their views solely through the use of hashtags rather than through direct exchanges with other users. For this reason, I create a discussion network among the participants on X (and also for Wikipedia), enabled us to observe the topics under discussion and to extract a comprehensive map of user interactions across different thematic areas.

For both databases, stopwords were used to clean the database and Bigram were use for the topics extraction. A bigram is a sequence of two adjacent elements from a string of tokens, which are typically letters, syllables, or words. A bigram is a sequence of two adjacent elements from a string of tokens, which are typically letters, syllables, or words. A bigram is an n-gram for $n = 2$ [24]. The frequency distribution of every bigram in a string is commonly used for simple statistical analysis of text in many applications, including computational linguistics, cryptography, speech recognition, and so on. In this case, the most important used words (given by the higher amount of connection between the words in the same sentence) were selected as nodes and linked (edges) with the second most used words in the same sentence.

The bigram probability is calculated using the following formula:

$$P(w_i | w_{i-1}) = \frac{C(w_{i-1}, w_i)}{C(w_{i-1})}$$

where:

- $P(w_i | w_{i-1})$ represents the conditional probability of word w_i given the previous word w_{i-1} ;
- $C(w_{i-1}, w_i)$ is the count of occurrences of the word pair (w_{i-1}, w_i) ;
- $C(w_{i-1})$ is the count of occurrences of the word w_{i-1} .

The probability of an entire sequence of words w_1, w_2, \dots, w_n can be calculated as:

$$P(w_1, w_2, \dots, w_n) = P(w_1) \prod_{i=2}^n P(w_i | w_{i-1})$$

As ending result, a networks of topics is created. The X database was divided between pro-war and anti-war to highlight different opinions and differences in the community. To avoid personal bias problems, Amazon Mechanical Turk (MTurk) was used to profile the two communities.

Each user was asked whether the sentence was Pro-war, Against-war , neutral or none at all. Also, a 5 minutes of time was left for each sentence so as not to rush users and to have better quality results. The same sentence was read and

Table 5. Distribution of tag combinations by stance.

Tag Combination	Pro-war	Against-war	Average Total
1 Tag	23.46%	18.52%	20.99%
2 Tags	59.26%	55.56%	57.41%
3 Tags	17.28%	25.93%	21.60%

<https://doi.org/10.1371/journal.pcsy.0000087.t005>

evaluated by at least 3 people. If the tweet being analysed shares at least two-thirds of its tags with a given tag set, the sentence is labelled accordingly. Table 5 presents the distribution of tags expressed as percentages. Also, to avoid political bias only individuals who do not reside in countries such as Russia, Ukraine, or Belarus were selected as MTurk workers. The results in the table show that 78.4% of the sample received a tag, while the remaining portion, consisting of items with at least three versions and conflicting opinions, was excluded.

With regard to the technical aspects of sentiment-analysis and emotion-detection models, I used the most suitable models for the dataset at hand, which was almost entirely in English and contained a large number of reviews. The choice was made to optimise both performance and accuracy on this specific data. For sentiment analysis I used the model available at [nlptown/bert-base-multilingual-uncased-sentiment](https://huggingface.co/nlptown/bert-base-multilingual-uncased-sentiment), which is pretrained on approximately 630,000 human-labeled product reviews across six languages. On a held-out test set of 30,000 reviews, this model achieves 95% off-by-1 accuracy, a standard reliability metric for ordinal sentiment scales.

For emotion detection, I used [distilbert-base-uncased-emotion](https://huggingface.co/distilbert-base-uncased-emotion), finetuned on a human-annotated Twitter emotion dataset. This model reports 93.8% accuracy and 93.79 F1, performing on par with or above heavier architectures (e.g. BERT, RoBERTa) on the same benchmark.

Because the models were applied as validated checkpoints and not re-trained on this corpus, I did not compute new confusion matrices. Instead, I rely on their published validation against human-labeled data. A methodological limitation remains the potential for domain shift in war-related discourse (e.g. sarcasm, euphemisms, culturally specific affective expressions), which is acknowledged and discussed in the limitations section.

A possible limitation of the sentiment-analysis and emotion-detection components of this study lies in the absence of account-level demographic information, such as gender or education. Previous research [25,26] has shown that gender-specific communication patterns can influence sentiment result in online discussions. However, neither the scraping procedures nor the X API provide reliable or ethically accessible demographic identifiers. Moreover, attempting to infer such attributes from usernames, profile pictures, or linguistic cues would introduce methodological biases and ethical concerns. For these reasons, gender-specific analyses could not be implemented in the present study.

In conclusion, in order to analyse the interaction between users of both X and Wikipedia, sentiment analysis, emotion detection (given the emphasis a war can bring) through BERT, and Shannon entropy were analysed using NLP techniques to observe whether there are differences between conflictual interaction and cooperation between simple and complex language.

Supporting information

S1 File. This work have its own limitations, such as the lack of Russian activity on X for a few days after the event, which restricts the comparison between the two platforms. A larger dataset could have revealed temporal fluctuations in sentiment and entropy across different events. Another key limitation is the potential presence of bots in the data.

(RAR)

Acknowledgments

This work was made possible thanks to the visiting period spent at University College Dublin (UCD) in 2023 with Taha Yasseri, supported by the University of Catania – Department of Physics “Ettore Majorana”, where I was enrolled during my doctoral studies, under the supervision of Professor Giovanni Giuffrida.

Author contributions

Conceptualization: Andrea Russo.

Data curation: Andrea Russo.

Formal analysis: Andrea Russo.

Investigation: Andrea Russo.

Resources: Andrea Russo.

References

1. Yasseri T, Menczer F. Can crowdsourcing rescue the social marketplace of ideas?. *Commun ACM*. 2023;66(9):42–5. <https://doi.org/10.1145/3578645>
2. Mestyán M, Yasseri T, Kertész J. Early prediction of movie box office success based on Wikipedia activity big data. *PLoS One*. 2013;8(8):e71226. <https://doi.org/10.1371/journal.pone.0071226> PMID: 23990938
3. Yasseri T, Sumi R, Rung A, Kornai A, Kertész J. Dynamics of conflicts in Wikipedia. *PLoS One*. 2012;7(6):e38869. <https://doi.org/10.1371/journal.pone.0038869> PMID: 22745683
4. Sumi R, Yasseri T, Rung A, Kornai A, Kertész J. Edit Wars in Wikipedia. In: 2011 IEEE Third Int'l Conference on Privacy, Security, Risk and Trust and 2011 IEEE Third Int'l Conference on Social Computing. 2011. p. 724–7. <https://doi.org/10.1109/passat/socialcom.2011.47>
5. Tsvetkova M, Yasseri T, Meyer ET, Pickering JB, Engen V, Walland P, et al. Understanding human-machine networks. *ACM Comput Surv*. 2017;50(1):1–35. <https://doi.org/10.1145/3039868>
6. Cinelli M, Morales GDF, Galeazzi A, Quattrocioni W, Starnini M. Echo chambers on social media: a comparative analysis. *arXiv preprint 2020*. <https://arxiv.org/abs/2004.09603>
7. Kaluža J. Far-reaching effects of the filter bubble, the most notorious metaphor in media studies. *AI & Soc*. 2022;38(4):1391–3. <https://doi.org/10.1007/s00146-022-01399-x>
8. Wilkinson DM, Huberman BA. Cooperation and quality in wikipedia. In: *Proceedings of the 2007 International Symposium on Wikis*. 2007. p. 157–64. <https://doi.org/10.1145/1296951.1296968>
9. Fengjun L, Zhengkui L, Na Z. Exploring the influence factors of collaborative conflicts in online knowledge community: an empirical research on Wikipedia. *Science Research Management*. 2019;40(3):153.
10. Viégas FB, Wattenberg M, Dave K. Studying cooperation and conflict between authors with history flow visualizations. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*; 2004. p. 575–82.
11. Yasseri T, Kornai A, Kertész J. A practical approach to language complexity: a Wikipedia case study. *PLoS One*. 2012;7(11):e48386. <https://doi.org/10.1371/journal.pone.0048386> PMID: 23189130
12. O’Leary DE, Storey VC. A Google–Wikipedia–Twitter model as a leading indicator of the numbers of coronavirus deaths. *Intelligent Sys in Account*. 2020;27(3):151–8. <https://doi.org/10.1002/isaf.1482>
13. Banda JM, Tekumalla R, Wang G, Yu J, Liu T, Ding Y, et al. A large-scale COVID-19 Twitter chatter dataset for open scientific research—an international collaboration. *Epidemiologia (Basel)*. 2021;2(3):315–24. <https://doi.org/10.3390/epidemiologia2030024> PMID: 36417228
14. Ishikawa S, Arakawa Y, Tagashira S, Fukuda A. Hot topic detection in local areas using Twitter and Wikipedia. In: *ARCS 2012. IEEE*; 2012. p. 1–5.
15. Tan L, Zhang H, Clarke C, Smucker M. Lexical comparison between Wikipedia and Twitter corpora by using word embeddings. In: *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*. 2015. p. 657–61. <https://doi.org/10.3115/v1/p15-2108>
16. Massa P. Social networks of Wikipedia. In: *Proceedings of the 22nd ACM conference on Hypertext and hypermedia*. 2011. p. 221–30. <https://doi.org/10.1145/1995966.1995996>
17. Zollo F, Novak PK, Del Vicario M, Bessi A, Mozetič I, Scala A, et al. Emotional dynamics in the age of misinformation. *PLoS One*. 2015;10(9):e0138740. <https://doi.org/10.1371/journal.pone.0138740> PMID: 26422473
18. Ferrara E, Yang Z. Measuring emotional contagion in social media. *PLoS One*. 2015;10(11):e0142390. <https://doi.org/10.1371/journal.pone.0142390> PMID: 26544688

19. Charquero-Ballester M, Walter JG, Rybner AS, Nissen IA, Enevoldsen KC, Bechmann A. Emotions on Twitter as crisis imprint in high-trust societies: do ambient affiliations affect emotional expression during the pandemic?. *PLoS One*. 2024;19(3):e0296801. <https://doi.org/10.1371/journal.pone.0296801> PMID: 38442085
20. Huszár F, Ktena SI, O'Brien C, Belli L, Schlaikjer A, Hardt M. Algorithmic amplification of politics on Twitter. *Proc Natl Acad Sci U S A*. 2022;119(1):e2025334119. <https://doi.org/10.1073/pnas.2025334119> PMID: 34934011
21. Bartley N, Abeliuk A, Ferrara E, Lerman K. Auditing algorithmic bias on twitter. In: *Proceedings of the 13th ACM Web Science Conference 2021*; 2021. p. 65–73.
22. Bovet A, Makse HA. Influence of fake news in Twitter during the 2016 US presidential election. *Nat Commun*. 2019;10(1):7. <https://doi.org/10.1038/s41467-018-07761-2> PMID: 30602729
23. Rudas C, Surányi O, Yasseri T, Török J. Understanding and coping with extremism in an online collaborative environment: a data-driven modeling. *PLoS One*. 2017;12(3):e0173561. <https://doi.org/10.1371/journal.pone.0173561> PMID: 28323867
24. Russo A, Miracula V, Picone A. Topics evolution through multilayer networks; analysing 2M tweets from 2022 Qatar FIFA World Cup. *arXiv preprint* 2024. <https://doi.org/arXiv:240112228>
25. Hudders L, De Jans S. Gender effects in influencer marketing: an experimental study on the efficacy of endorsements by same- vs. other-gender social media influencers on Instagram. *International Journal of Advertising*. 2021;41(1):128–49. <https://doi.org/10.1080/02650487.2021.1997455>
26. Thakur N, Cui S, Khanna K, Knieling V, Duggal YN, Shao M. Investigation of the gender-specific discourse about online learning during COVID-19 on Twitter using sentiment analysis, subjectivity analysis, and toxicity analysis. *Computers*. 2023;12(11):221. <https://doi.org/10.3390/computers12110221>