

Modeling Linguistic Imprints of War Propaganda in a Russian Wikipedia Fork: A Comparative Analysis with the Original Wikipedia

Anastasiia Vestel

Saarland University

Saarbruecken, Germany

anastasiia.vestel@uni-saarland.de

Stefania Degaetano-Ortlieb

Saarland University

Saarbruecken, Germany

s.degaetano@mx.uni-saarland.de

Abstract

Although Wikipedia aspires to provide neutral information, alternative versions can be used for political manipulation. This paper analyzes how narratives about the Russo-Ukrainian War are linguistically reframed in a Russian Wikipedia Fork compared to the original Russian Wikipedia. Using Kullback-Leibler Divergence on a corpus of war-related edits in more than 13,000 articles, we identify key differences between the two versions. While the original Wikipedia features Ukrainian references and administrative details, direct war terminology, and Ukraine’s territorial designation, governance, and statehood, RWFork replaces or removes these elements, emphasizing reassignment of Ukrainian territories to Russia, favoring euphemistic war language, renaming locations, and recognizing Russia-backed DPR and LPR¹. These patterns closely align RWFork with demobilizational strategies observed in pro-Kremlin media.

1 Introduction and Related Work

Language, being a powerful tool, can be manipulated by malicious actors to shape public opinion. In this study, we examine how pro-Kremlin propaganda about the Russo-Ukrainian War is linguistically framed in a Russian Wikipedia Fork (RWFork) compared to the original Russian Wikipedia. RWFork, or Ruwiki, is a copy of Russian Wikipedia created in June 2023 and revised to comply with Russia’s legislation. It was launched by Vladimir Medeyko, the leader of Russian Wikipedia editors, who aspired to create a Kremlin-aligned alternative to the original Wikipedia (Cohen, 2023). Despite his denial about working for the Russian government, RWFork has

¹DPR (Donetsk People’s Republic) and LPR (Luhansk People’s Republic) are self-proclaimed separatist entities in Donbas that emerged in 2014, supported and recognized by Russia, which are internationally viewed as Ukrainian territory under Russian occupation.

been described as sponsored by the state and selectively edited to remove content unfavorable to the Kremlin (Roscoe, 2024).

Previous work on manipulative narratives in the context of the Russo-Ukrainian War has shown their evolving nature and adaptation to external events (Gerard et al., 2025a,b; Solopova et al., 2023). Therefore, RWFork presents a compelling case study for exploring the development of Russian propaganda by comparing two versions of Wikipedia separated by an important historical event: the start of Russia’s full-scale invasion of Ukraine.

As our method, we employ Kullback-Leibler Divergence (KLD; Kullback and Leibler, 1951), which, as opposed to transformer-based methods, offers an interpretable way to identify distinctive linguistic features. KLD has been established as a reliable method for detecting language variation and change, such as in studies by Bochkarev et al. (2014); Degaetano-Ortlieb and Teich (2019); Fankhauser et al. (2014); Hughes et al. (2012); Klingenstein et al. (2014); among others. Moreover, we have previously demonstrated that KLD is suitable for detecting divergences in propaganda strategies between traditional and social media covering the Russo-Ukrainian War (Vestel and Degaetano-Ortlieb, 2025). This study extends prior work by applying KLD to edits in RWFork, thus adding another text type (Wikipedia) and a diachronic dimension (the difference between the original and changed versions) to the analysis.

Propaganda studies on the Russo-Ukrainian War have mostly focused on media analysis (Akhynko et al., 2025; Hein, 2023; Vanetik et al., 2023, etc.). For instance, Alyukov et al. (2025) examined the differences between state-controlled traditional media (such as press and TV) and social media, showing that the former rely on demobilization by normalizing the war, while the latter have a more mobilizational character. A prominent

example of the normalization frame in traditional media is replacing direct war-related terminology with euphemisms such as *special military operation* (Alyukov et al., 2022; Vestel and Degaetano-Ortlieb, 2025). A similar trend was found in other media types, specifically in state-affiliated outlets as opposed to independent ones (Park et al., 2022) and in pro-Russian vs. pro-Ukrainian Telegram channels (Ustyianovych and Barbosa, 2024). More broadly, Gerard et al. (2025a) have demonstrated that Russia’s Telegram discourse seems to be detached from the realities of war, turning attention towards international responses and economic repercussions, which also aligns with a demobilizational strategy. Likewise, Russian bloggers more frequently focus on international actions instead of military developments, which is more characteristic of Ukrainian bloggers (Gerard et al., 2025b).

Wikipedia presents another text type, different from media: while aiming to be a neutral and objective source of information, it can also serve as a tool for knowledge manipulation, specifically in its alternative versions such as RWFork (Trokhymovych et al., 2025). Similarly to the above-mentioned studies on media propaganda, Trokhymovych et al. (2025) have shown that direct words like *war* and *invasion* are removed from RWFork. Moreover, RWFork replaces names of occupied Ukrainian regions with Kremlin-aligned terms and often removes Ukrainian references and administrative details (Trokhymovych et al., 2025). Both of these types of changes can be attributed to the normalization frame: the former aims to divert the readers’ attention from the war, whereas the latter tries to convince them that certain Ukrainian territories rightfully belong to Russia. Thus, RWFork seems to employ a demobilizational propaganda strategy, making it more similar to traditional media in contrast to social media.

The current study will test this hypothesis and not only investigate the divergences between the original and changed versions of Wikipedia but also compare the results with prior work on media analysis.

2 Data and Methodology

Our study uses Trokhymovych et al.’s (2025) RWFork dataset, containing edits from 1.9M page titles between May and September 2023. We filtered the data by selecting only the articles from the three categories of changes involving knowledge

manipulation linked to the 2022 Russian invasion of Ukraine, namely Territorial Claims Dispute, Terminology Changes Ukraine, and Sanctions Edit Adjustments (the category labels are assigned to each article; see the original paper for details). The resulting data comprised 13,048 articles (100,738 sentences in total). However, at first, we discovered that our data contained duplicated template sentences, requiring further pre-processing (refer to Appendix A for details). After removing them, the total number of sentences decreased to 92,629. Appendix B contains detailed statistics on the number of sentences, tokens, and analyzed words for each version before and after removing duplicates.

KLD is used to measure how much the language of RWFork diverges from Russian Wikipedia by quantifying the extra information needed to represent one probability distribution with the other, thereby identifying key linguistic features (e.g., words) that contribute to the linguistic differences. For instance, to calculate the KLD of the language of the *original* Wikipedia (O), given that of the *changed* RWFork (C), we use this formula:

$$D(O \parallel C) = \sum_i p(\text{feature}_i \mid O) \log_2 \frac{p(\text{feature}_i \mid O)}{p(\text{feature}_i \mid C)}$$

We applied KLD (Degaetano-Ortlieb and Teich, 2019) to the RWFork corpus, limiting the scope to the lemmas of nouns, proper nouns, verbs, adjectives, and adverbs, since these parts of speech (POS) are most likely to account for substantial differences in content words between the two versions of Wikipedia. The results, described in the next section, are based on the point-wise KLD, which measures the contribution of each feature to the overall KLD — high contribution reflects *distinctive* features. As opposed to other methods for detecting linguistic variation, such as the Jensen-Shannon Divergence, KLD is asymmetric, which makes it well-suited for capturing differences from distinct perspectives (Degaetano-Ortlieb and Teich, 2019).

Finally, we performed qualitative analysis by looking at the 50 most distinctive words for each version. In particular, we manually annotated each word with a topic it refers to. We drew inspiration from Trokhymovych et al.’s (2025) taxonomy of changes to classify the words according to topics, but the categories are more specific in our case, allowing for a more fine-grained analysis (see Appendix C for details). To find out the context around each word, we performed queries on the dataset us-

Topic	Original	Changed
DPR/LPR recognition	0.00%	10.00%
non-recognition of DPR/LPR	7.41%	0.00%
occupation/annexation (direct)	9.26%	0.00%
occupation/annexation (euphemism)	0.00%	4.00%
other	3.70%	10.00%
political figures	3.70%	0.00%
propaganda	1.85%	0.00%
renamed locations	0.00%	20.00%
sanctions	0.00%	8.00%
territorial designation & governance (Russia)	0.00%	34.00%
territorial designation & governance (Ukraine)	18.52%	0.00%
Ukraine's statehood	12.96%	0.00%
Ukrainian references & administrative details	22.22%	0.00%
war (direct)	20.37%	0.00%
war (euphemism)	0.00%	14.00%

Table 1: Percentages of words with the highest KLD values for the original and changed versions of Wikipedia grouped by topic.

ing Sketch Engine² as our corpus exploration tool. We used the parallel concordance feature, where the original and changed versions were aligned on the sentence level³, and examined the most frequent collocates for each word to inform our topic classification. In cases when a word could refer to more than one topic, it would receive multiple category labels.

3 Results

3.1 Overall Divergence

The overall KLD for the original Russian Wikipedia is close to 0.684, whereas for RWFork, the divergence is around 0.415, meaning that 0.269 additional bits are required to encode the former corpus with the latter. An explanation for this could be that RWFork is characterized by a more limited vocabulary, given that it contains many duplicated sentences (even after removing most of them, there were still a few template phrases left, as will be discussed in the next subsection).

3.2 Divergences between the Russian Wikipedia and RWFork

Figure 1 shows the most distinctive words for the Russian Wikipedia and RWFork (a list of these words, their translation into English, and their probabilities can be found in Appendix D). Additionally, Table 1 summarizes the percentage of words that belong to different topics in either version.

²<https://www.sketchengine.eu/>

³Note that only modified sentences were aligned; added and deleted sentences had empty entries in the Russian Wikipedia and RWFork, respectively.

About 22% of the words distinctive to the original version indicate Ukrainian references and administrative details in relation to certain locations occupied or claimed by Russia; they are often deleted from RWFork, implying that these territories are now part of Russia. Examples include *postal*, *telephone*, *code*, and *index*, the abbreviation *KOATUU* (which translates to "Classification of objects of the administrative-territorial system of Ukraine"), and links to the website of the Verkhovna Rada (Supreme Council), Ukraine's parliament (evidenced by the words *website*, *supreme*, *council*, and *Ukraine*).

The next most common category of words distinctive of the original Wikipedia is that of direct war terminology (around 20%), such as *invasion*, *war*, *aggression*, and *attack*, as well as their collocates in phrases like *full-scale invasion*, someone's *support* for the war, or *during* the war. These words and phrases are, for the most part, removed from RWFork or otherwise substituted with vague terms such as *conflict*, *military actions*, and *special military operation* (the words comprising these expressions are among the top KLD words for RWFork; see Figure 1 and Appendix D), as well as Russia's politics *in relation to* Ukraine. It is worth noting that 14% of the 50 most distinctive words for RWFork are war-related euphemisms or their collocates, making this category the third most common type of changed Wikipedia content. In a similar vein, direct words that refer to Russia's occupation of Ukrainian territories (e.g., *annexation*, *occupation*, *to occupy*, *to capture*) are distinctive for the Russian Wikipedia, comprising 9% of the top distinctive KLD words for this version; they are also mostly deleted or replaced with euphemisms such as *inclusion* and *entry* (into Russia).

Almost 19% of the words distinctive for the original version are those indicating the territorial designation and governance of Ukraine. For instance, the word for *council* (transliterated as *sovet*) and its collocate *rural council*, as a form of local government in Ukraine, are either removed or replaced with a corresponding form of local government in Russia (*okrug*). Another example is the word *autonomous*, which is among the top distinctive KLD words for the original Wikipedia and is often part of the phrase "Autonomous Republic of Crimea", as this is its official name according to the administrative division of Ukraine. After Russia annexed the peninsula in 2014, it renamed the territory to "Republic of Crimea", leaving out the word "au-



Figure 1: Top 50 words with the highest KLD values for the original (left) and changed versions (right). Size accounts for KLD values, and color indicates frequencies from higher (red) to lower (green).

onomous", which suggests its belonging to Russia. Therefore, this word was mostly deleted from the republic's name in RWFork.

A few words (approx. 13%) among those most distinctive for the Russian Wikipedia refer to the territorial statehood of Ukraine (such as *territorial*, *integrity*, *sovereignty*, and *independence*). They are mostly mentioned when talking about Russia's violation and disruption of Ukraine's statehood (hence, the words *violation* and *to disrupt* also have a high contribution to the language of the original version), and these phrases are almost always removed from RWFork.

More than one-third of the top distinctive KLD words for RWFork (34%) indicate the territorial designation and governance of Russia (including the so-called DPR and LPR, since they are backed by Russia). A lot of these words turned out to be a case of a few more template sentences that were duplicated in different articles and were not initially detected. These sentences state that before the year 2014 or 2022, a certain village was part of an occupied Ukrainian region; that since spring 2014, it has belonged to the DPR; or that since August 2022, it has been part of a city or municipal council of Kherson Oblast. These phrases contributed numerous words to the RWFork word cloud, such as *village*, *year*, *spring*, *August*, *municipal*, etc., implying that the specified territory now belongs to Russia. Such framing indicates territorial control over certain locations in the self-proclaimed DPR, as well as villages in Kherson Oblast captured by Russia during the full-scale invasion.

Furthermore, 20% of the top distinctive KLD words for RWFork are Ukrainian territories con-

trolled or claimed by Russia that have been renamed by the Russian government. This is often the case of administrative reorganization, when the borders of certain regions or districts are redrawn to their previous state during the Soviet regime, and their names are changed accordingly. A prominent example is *Krasnoarmeysk*, a former district in Donetsk Oblast partially occupied by Russia, which, according to the Ukrainian administrative division, is now part of Pokrovsk Raion (district). "Krasnoarmeysk" literally translates to "Red Army", and it was the Soviet name of the district (and its capital city) up until 2016, when it was renamed as Pokrovsk Raion by the Ukrainian government in the process of decommunization; in 2020, it was enlarged to include a few territories of other neighboring districts. By renaming certain locations in this way, the editors of RWFork aim to promote the narrative that they have inherently belonged to Russia.

Finally, the language of RWFork clearly indicates the recognition of the self-proclaimed DPR and LPR, as 10% of the most distinctive words for the changed version are related to those regions. These words are *LPR*, *Luhansk* (or *Lugansk*, if transliterated from Russian), *Donetsk*, *people's*, and *republic*. On the contrary, the original Wikipedia does not recognize these republics, which is evidenced by some words with high KLD values for this version. These include *self-proclaimed* and *separatist*, as well as *control* and *to be controlled* (in sentences like "[This territory] is controlled by the self-proclaimed DPR/LPR").

3.3 Discussion

By comparing our results with other studies on propaganda about the Russo-Ukrainian War, we have observed many similarities between the language of RWFork and that of traditional, state-affiliated, and pro-Russian media. The most remarkable example is that direct war-related terminology is often deleted (as was also noted by Trokhymovych et al., 2025) or replaced with vague euphemistic expressions, a pattern found by Alyukov et al. (2022); Park et al. (2022); Ustyianovych and Barbosa (2024); Vestel and Degaetano-Ortlieb (2025). This, and the fact that the words *war*, *invasion*, *aggression*, etc., are removed more often than they are substituted, points to a demobilizational strategy frequently employed by traditional media (Alyukov et al., 2025). Moreover, by renaming Ukrainian locations, deleting any references to Ukraine, and recognizing Russia-backed DPR and LPR, the authors of RWFork seek to convince the population that these territories have been historically part of Russia. These changes point to territorial control propaganda, which also reflects the normalization frame of the demobilizational approach, and the same trend was found to be distinctive for traditional media as opposed to social media in our previous study (Vestel and Degaetano-Ortlieb, 2025).

These findings reveal that studying propaganda in alternative versions of Wikipedia, such as RWFork, is equally important to media research. The encyclopedic style of RWFork gives it a certain degree of credibility as opposed to media sources, and the demobilizational strategies employed in it are a more subtle propaganda technique than mobilization. Therefore, ordinary users might rely on RWFork as an objective source of information without questioning its validity or suspecting possible manipulation. For example, while someone might not be searching for the news on purpose but stumbles upon an article about a certain town or village, they might see the location marked as "Russia" and not realize that this is Ukrainian territory occupied by Russia as a result of the invasion.

4 Conclusion and Future Work

By employing KLD to compare the two versions of Russian Wikipedia, we have identified key linguistic features that distinguish propagandistic rhetoric in war-related narratives. Specifically, the language of the original version is characterized by Ukrainian references and administrative details,

Ukraine's territorial designation, governance, and statehood, as well as direct terminology for the war and Russia's occupation or annexation of Ukrainian territories. In contrast, RWFork focuses on Russia's territorial designation and governance, renaming Ukrainian locations, euphemistic framing of the war, and recognizing the self-proclaimed DPR and LPR. In addition, we discovered that the language of RWFork aligns closely with pro-Kremlin propaganda spread in pro-Russian, state-affiliated, or traditional media, as opposed to the original version. This is evidenced by the efforts to downplay the war and the focus on territorial control, which point to a demobilizational strategy. Overall, our work sheds light on how propaganda is construed in political environments and enhances reproducible computational techniques for its detection.

In the future, we will expand our set of linguistic features (e.g., dependency relations) to investigate more subtle propagandistic framing. Furthermore, we aim to extend this study by applying other methods to the dataset, such as surprisal (Shannon, 1948), which models the (un)expectedness of words in particular contexts to capture more nuanced local linguistic changes. Finally, for a more fine-grained analysis and classification into different propaganda techniques, such as those proposed by Da San Martino et al. (2019), a comprehensive study of evaluative language is needed; this can be informed by word embeddings (Mikolov et al., 2013), which will allow us to model semantic divergences between the two versions of Wikipedia.

Acknowledgments

Funded by the European Union under grant 101119511. Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or European Research Executive Agency (REA). Neither the European Union nor the granting authority can be held responsible for them.

We sincerely thank Mykola Trokhymovych for his help with the RWFork dataset and Stefan Fischer for providing technical support.

References

Kateryna Akhynko, Oleksandr Kosovan, and Mykola Trokhymovych. 2025. *Hidden Persuasion: Detecting Manipulative Narratives on Social Media During the 2022 Russian Invasion of Ukraine*. In *Proceedings of the Fourth Ukrainian Natural Language Process-*

- ing Workshop (UNLP 2025), pages 194–202, Vienna, Austria (online). Association for Computational Linguistics.
- Maxim Alyukov, Maria Kunilovskaya, and Andrei Semenov. 2025. *Confuse and Normalise: Authoritarian Propaganda in a High-Choice Media Environment and Russia’s Invasion of Ukraine*. In Paul Goode, editor, *Russian Propaganda Today: Challenges, Effectiveness, and Resistance*, page in print. University of Michigan press, University of Manchester Press.
- Maxim Alyukov, Andrei Semenov, and Maria Kunilovskaya. 2022. Propaganda Setbacks and Appropriation of Anti-war language: “Special Military Operation” in Russian Mass Media and Social Networks (February-July 2022). Monitoring Report №1.
- Vladimir Bochkarev, Valery D. Solovyev, and Søren Wichmann. 2014. *Universals versus Historical Contingencies in Lexical Evolution*. *Journal of The Royal Society Interface*, 11(101):20140841.
- Noam Cohen. 2023. Russian Wikipedia’s Top Editor Leaves to Launch a Putin-Friendly Clone. *Bloomberg.com*.
- Giovanni Da San Martino, Seunghak Yu, Alberto Barrón-Cedeño, Rostislav Petrov, and Preslav Nakov. 2019. *Fine-Grained Analysis of Propaganda in News Article*. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5636–5646, Hong Kong, China. Association for Computational Linguistics.
- Stefania Degaetano-Ortlieb and Elke Teich. 2019. *Toward an Optimal Code for Communication: The Case of Scientific English*. *Corpus Linguistics and Linguistic Theory*, 18(1):175–207.
- Peter Fankhauser, Jörg Knappen, and Elke Teich. 2014. Exploring and Visualizing Variation in Language Resources. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC’14)*, pages 4125–4128, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Patrick Gerard, Svitlana Volkova, Louis Penafiel, Kristina Lerman, and Tim Weninger. 2025a. *Modeling Information Narrative Evolution on Telegram During the Russia-Ukraine War*. *Proceedings of the International AAAI Conference on Web and Social Media*, 19:602–614.
- Patrick Gerard, Tim Weninger, and Kristina Lerman. 2025b. *Fear and Loathing on the Frontline: Decoding the Language of Othering by Russia-Ukraine War Bloggers*. *Proceedings of the International AAAI Conference on Web and Social Media*, 19:615–635.
- Vitalij Hein. 2023. *Propaganda Detection in Russian and American News Coverage about the War in Ukraine through Text Classification*. Thesis, Technische Universität Wien.
- James M. Hughes, Nicholas J. Foti, David C. Krakauer, and Daniel N. Rockmore. 2012. *Quantitative Patterns of Stylistic Influence in the Evolution of Literature*. *Proceedings of the National Academy of Sciences*, 109(20):7682–7686.
- Sara Klingenstein, Tim Hitchcock, and Simon DeDeo. 2014. *The Civilizing Process in London’s Old Bailey*. *Proceedings of the National Academy of Sciences of the United States of America*, 111(26):9419–9424.
- Solomon Kullback and Richard A. Leibler. 1951. *On Information and Sufficiency*. *The Annals of Mathematical Statistics*, 22(1):79–86.
- Tomás Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient Estimation of Word Representations in Vector Space. In *1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2-4, 2013, Workshop Track Proceedings*.
- Chan Young Park, Julia Mendelsohn, Anjalie Field, and Yulia Tsvetkov. 2022. *Challenges and Opportunities in Information Manipulation Detection: An Examination of Wartime Russian Media*. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 5209–5235, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Jules Roscoe. 2024. Russia Clones Wikipedia, Censors It, Bans Original. <https://www.404media.co/russia-clones-wikipedia-censors-it-bans-original/>.
- Claude E. Shannon. 1948. *A Mathematical Theory of Communication*. *Bell System Technical Journal*, 27(3):379–423.
- Veronika Solopova, Christoph Benz Müller, and Tim Landgraf. 2023. *The Evolution of Pro-Kremlin Propaganda From a Machine Learning and Linguistics Perspective*. In *Proceedings of the Second Ukrainian Natural Language Processing Workshop (UNLP)*, pages 40–48, Dubrovnik, Croatia. Association for Computational Linguistics.
- Mykola Trokhymovych, Oleksandr Kosovan, Nathan Forrester, Pablo Aragón, Diego Saez-Trumper, and Ricardo Baeza-Yates. 2025. *Characterizing Knowledge Manipulation in a Russian Wikipedia Fork*. *Proceedings of the International AAAI Conference on Web and Social Media*, 19:1924–1936.
- Taras Ustyianovych and Denilson Barbosa. 2024. Instant Messaging Platforms News Multi-Task Classification for Stance, Sentiment, and Discrimination Detection. In *Proceedings of the Third Ukrainian Natural Language Processing Workshop (UNLP) @ LREC-COLING 2024*, pages 30–40, Torino, Italia. ELRA and ICCL.
- Natalia Vanetik, Marina Litvak, Egor Reviakin, and Margarita Tyamanova. 2023. *Propaganda Detection in Russian Telegram Posts in the Scope of the Russian Invasion of Ukraine*. In *Proceedings of the Conference Recent Advances in Natural Language*

Anastasiia Vestel and Stefania Degaetano-Ortlieb. 2025. [From War to Special Military Operation: Interpretable Detection of Linguistic Propaganda Framing in Russian Media](#). *Workshop Proceedings of the 19th International AAAI Conference on Web and Social Media*, 2025:50.

A Duplicated sentences

During the initial analysis, we observed that 52% of the top KLD words for RWFork came from only three template sentences used throughout various articles about different Ukrainian locations occupied or claimed by Russia. These sentences exhibited changes in territorial designation and governance, asserting that, according to an agreement between a certain Ukrainian region and the Russian Federation, the specified location belonged to Russia; that the Ukrainian authorities disputed this; and that Ukraine exercised de facto control over this territory. Examples of words that stem from these sentences and have a high contribution to the language of RWFork are *agreement*, *to dispute*, and *de facto* (see Figure 2). A possible explanation for the repetition of these template texts is that a significant share of RWFork’s editorial work may have been carried out by paid staff (Trokhymovych et al., 2025). If this is the case, the RWFork contributors might have received certain guidelines that prescribed inserting these sentences in multiple articles to unify RWFork’s content across different pages. This can be regarded as a specific type of information manipulation: for instance, a study by Da San Martino et al. (2019) showed that repetition was the third most common strategy within a fine-grained taxonomy of 18 propaganda techniques.

Since the duplicated sentences impacted our KLD results to such a great extent, we decided to remove them and rerun the experiment. The motivation behind this was the fact that we wanted to analyze how propaganda is framed linguistically compared to the original Wikipedia version, and the duplicated content accounted mostly for extralinguistic factors, therefore overshadowing more meaningful features that might point to certain manipulative strategies. After deleting the duplicates, we saw that 80% of the top 50 words distinctive for the original Wikipedia were the same as in the initial setup of the experiment, but the overlap between the 50 words with the highest KLD values

for RWFork with and without the repeated template sentences was only 58%.

B Descriptive statistics of the data

Table 2 shows the number of sentences, tokens, and POS-filtered target words used in our analysis for each version and in total, both before and after removing the duplicated template sentences discussed in Appendix A. The number of articles reported in Section 2 is the same for both versions.

C Classification of topics

Below is a detailed taxonomy of topics, expanded from the three initial categories of changes related to the Russo-Ukrainian War, which were proposed by Trokhymovych et al. (2025). Apart from making the three clusters more specific, we add more categories (under "Other"). In addition, we differentiate between a pro-Russian and a pro-Ukrainian stance by separating direct terminology from euphemisms, as well as discriminating between the language that does or does not recognize the DPR and LPR and the phrasing that supports Russia’s or Ukraine’s claim on a territory.

- Territorial Claims Dispute
 - DPR/LPR recognition
 - non-recognition of DPR/LPR
 - renamed locations
 - territorial designation and governance (Russia)
 - territorial designation and governance (Ukraine)
 - Ukrainian references and administrative details
- Terminology Changes Ukraine
 - occupation/annexation (direct)
 - occupation/annexation (euphemism)
 - war (direct)
 - war (euphemism)
- Sanctions Edit Adjustments
 - sanctions
- Other
 - political figures
 - propaganda
 - Ukraine’s statehood
 - other



Figure 2: Top 50 words with the highest KLD values from the original analysis for the Russian Wikipedia (left) and RWForK (right). Size accounts for KLD values, and color indicates frequencies from higher (red) to lower (green).

#	Before deduplication			After deduplication		
	Original	Changed	Total	Original	Changed	Total
Sentences	59,648	41,090	100,738	59,648	32,981	92,629
Tokens	859,234	632,051	1,491,285	859,234	477,274	1,336,508
Words analyzed	517,696	409,322	927,018	517,696	298,291	815,987

Table 2: Descriptive statistics of the data used for the analysis.

D Top 50 KLD words

Table 3 includes a list of 50 words with the highest KLD values, their translation into English, and their probabilities for the original and changed versions of Wikipedia. Note that to translate the names of Ukrainian locations, we did not use their official translations from Ukrainian into English but transliterated them from Russian into Latin script, since many of them are the ones that have been renamed (see Section 3.2); others, like *Zaporozh'ye* instead of *Zaporizhzhia*, have also been transliterated and not translated for consistency. However, *Luhansk* in the word list for the original Russian Wikipedia was translated in this way because this word was originally in Ukrainian (it appeared as part of an address), while its Russian-to-English transliteration, *Lugansk*, can be seen in the word list for RWForK, since in this version it was spelled in Russian. At the same time, most of the locations are adjectives used in the names of regions (e.g., *Kherson Oblast*), so we transliterated the city or town their names come from. The repeated word *Chaplynka* (name of a region) is a result of a parsing error when different word forms were not recognized as the same lemma. Finally, abbreviations that have not been mentioned in the paper are explained in parentheses.

words	translation	kld1	kld2	p1	p2
вторжение	invasion	0.08660169	-0.000542759	0.011774726	7.42601E-05
код	code	0.040392951	-0.001003039	0.007576064	0.000188129
коатуу	koatuu	0.032426735	-0.000178471	0.00432048	2.37791E-05
украина	ukraine	0.031051865	-0.018145495	0.040063383	0.023411474
индекс	index	0.020182594	-0.000198151	0.003025711	2.97062E-05
сайт	website	0.020110603	-0.001344424	0.005152738	0.000344468
рада	council	0.018103583	-0.001330889	0.004807351	0.000353414
верховный	supreme	0.016541175	-0.001630482	0.00494846	0.000487775
телефонный	telephone	0.015244526	-0.000608512	0.003280607	0.000130951
россия	russia	0.013438068	-0.007987493	0.017905223	0.010642738
почтовый	postal	0.01141654	-0.000860441	0.003060813	0.000230687
территориальный	territorial	0.010842433	-0.002122099	0.004607674	0.000901822
аннексия	annexation	0.009574253	-3.04384E-05	0.001153924	3.66855E-06
война	war	0.009325456	-0.002248677	0.004544352	0.001095794
самопровозглашённый	self-proclaimed	0.008389039	-0.000312659	0.00176766	6.58805E-05
совет	soviet	0.008257453	-0.003712976	0.007160961	0.003219936
деление	division	0.006969798	-0.000601835	0.001972392	0.000170314
луганська	luhansk	0.006099984	-1.9393E-05	0.000735193	2.33732E-06
российский	russian	0.005500208	-0.003841097	0.010618824	0.007415708
целостность	integrity	0.005355526	-0.000634959	0.001740909	0.000206404
украинской	ukrainian	0.005183336	-0.001723559	0.003263059	0.001085025
поддержка	support	0.004810887	-0.001659934	0.003133767	0.001081266
агрессия	aggression	0.004757816	-0.00040701	0.001341302	0.000114742
административный	administrative	0.004329221	-0.001950613	0.003763952	0.001695921
сельский	rural	0.00406735	-0.002547075	0.006023492	0.003772059
оккупировать	occupy	0.00378714	-0.000346861	0.001098143	0.000100578
население	population	0.003690449	-0.001834607	0.003659979	0.00181946
улице	street	0.003683599	-6.92844E-05	0.000642588	1.20864E-05
нарушение	violation	0.003490644	-0.000517941	0.00126811	0.000188162
окупация	occupation	0.003377183	-0.000146988	0.000746826	3.25047E-05
контроль	control	0.003344609	-0.001133892	0.002143218	0.000726596
российско	russo-	0.002988045	-0.001009512	0.001907661	0.000648847
классификация	classification	0.002898066	-0.000105284	0.000605943	2.20133E-05
время	time	0.002808515	-0.001581546	0.003899955	0.001908993
область	oblast	0.002346483	-0.002118394	0.01590522	0.014359159
автономный	autonomous	0.002212086	-0.001023394	0.001989203	0.00092028
подрывать	disrupt	0.0021831	-0.000176871	0.000602133	4.87838E-05
суверенитет	sovereignty	0.002001847	-0.000406664	0.000870586	0.000176855
обл	obl (oblast)	0.001888448	-0.001187279	0.002820506	0.001773269
так	so	0.001827927	-0.000800015	0.001533354	0.000671091
путин	putin	0.001819609	-0.001051689	0.002300623	0.001329704
независимость	independence	0.001639769	-0.000499208	0.000955699	0.000290951
контролировать	be controlled	0.001490477	-9.70816E-05	0.000378252	2.46373E-05
устройство	structure	0.001446351	-0.000227786	0.000542382	8.54199E-05
захватить	capture	0.001418415	-0.000274318	0.000598397	0.000115728
пропаганда	propaganda	0.001383671	-0.000224217	0.000526945	8.53708E-05
сепаратист	separatist	0.001332055	-0.000120349	0.00038406	3.46992E-05
полномасштабный	full-scale	0.001315102	-0.000150678	0.000420747	4.82071E-05
нападение	attack	0.001310835	-0.000208434	0.000494129	7.85708E-05
рф	rf (russian federation)	0.001277267	-0.000861606	0.002248867	0.001517017

words	translation	kld1	kld2	p1	p2
народный	people's	-0.006554873	0.038029752	0.002584231	0.014993069
республика	republic	-0.011293984	0.032744543	0.007354291	0.021322227
военный	military	-0.006373787	0.030521043	0.002820782	0.013507388
действие	action	-0.005416492	0.027148564	0.002329228	0.011674566
село	village	-0.012863659	0.024819316	0.013566939	0.026176233
год	year	-0.012912389	0.017643704	0.028668868	0.039173622
район	raion	-0.010701834	0.016946686	0.016137995	0.025555017
входить	be part of	-0.003767267	0.01007444	0.002654669	0.007099128
округ	okrug	-0.001425247	0.009956661	0.005082029	0.003550307
состав	composition	-0.003387915	0.006511743	0.003594048	0.00690794
херсонский	kherson	-0.003165557	0.006157423	0.003297911	0.006414869
донецкий	donetsk	-0.004173868	0.005775763	0.008906591	0.012324867
данным	data	-0.001425812	0.005676612	0.000715322	0.002847923
посёлок	village	-0.002631359	0.004569797	0.00330439	0.005738629
лнр	lpr	-0.001559755	0.004148277	0.001105274	0.002393554
запорожский	zaporozh'ye	-0.002244964	0.004000972	0.002692907	0.004799295
муниципальный	municipal	-0.000874695	0.003969679	0.000400837	0.001819142
городской	city	-0.002108014	0.003279932	0.003305228	0.00514272
луганский	lugansk	-0.002294771	0.003054	0.005565111	0.007406337
санкционный	sanctioned	-0.00167688	0.002443285	0.003087931	0.004499247
список	list	-0.00168797	0.00237826	0.003412674	0.004808279
начало	start	-0.001418093	0.00225028	0.002128781	0.003378025
тип	type	-0.001093008	0.002243565	0.001053681	0.002162695
вольнянском	volnyansk	-4.59238E-05	0.002031891	8.39951E-06	0.000371635
отношение	relation	-0.000930081	0.001721122	0.001047481	0.001938371
красноармейский	krasnoarmeysk	-0.000131614	0.001709402	3.55801E-05	0.000462113
мо	mo (municipal okrug)	-0.000115489	0.001702974	2.97482E-05	0.000438659
артёмовский	artemovsk	-0.000190333	0.001667124	6.07942E-05	0.000532497
км	km	-0.000998368	0.00160618	0.001455371	0.002341407
поселение	settlement	-0.001016529	0.001601508	0.001550107	0.002442143
специальный	special	-0.00050051	0.00151948	0.000312408	0.000948427
относится	belong	-0.000893753	0.001452661	0.001275423	0.002073009
краснолиманском	krasnyy liman	-1.51591E-05	0.001398687	2.32226E-06	0.000214268
август	august	-0.000725937	0.001321208	0.000840261	0.001529278
операция	operation	-0.000607059	0.001311385	0.000546318	0.00118017
численность	number	-0.000349115	0.00122901	0.000192273	0.000676871
присоединение	inclusion	-0.000545074	0.001178221	0.000490137	0.00105947
февраль	february	-0.000995856	0.001170766	0.004265949	0.005015209
александровский	aleksandrovka	-0.000475038	0.001097101	0.000393379	0.00090851
воладарский	volodarskoye	-9.98626E-05	0.001042632	2.9509E-05	0.000380893
чапльинский	chaplinsk	-1.89641E-06	0.001034915	2.08579E-07	0.000113827
сша	usa	-0.000751359	0.000877845	0.000347352	0.003910855
гуляйпольском	gulyaypole	-5.5183E-05	0.000863695	1.39062E-05	0.000217653
вхождение	entry	-0.000160475	0.000859995	6.6258E-05	0.00035508
куйбышевский	kuybyshevo	-0.000119403	0.000818604	4.29919E-05	0.000234745
весна	spring	-0.000341452	0.000808923	0.000275199	0.000650352
чапльинском	chaplinsk	-1.29288E-05	0.000795045	2.17503E-06	0.00013392
город	city	-0.00067337	0.000763929	0.003699068	0.004196538
конфликт	conflict	-0.000531337	0.00074783	0.001077582	0.001516642
значение	significance	-0.000396116	0.000733817	0.000445326	0.00082498

Table 3: The 50 most distinctive words for the original (left) and changed versions (right), together with their English translations and respective probabilities ("p"); "kld1" and "p1" are used for the Russian Wikipedia, while "kld2" and "p2" represent values for RWFork.