



This is a repository copy of *Comparative analysis of engagement, themes, and causality of Ukraine-related debunks and disinformation*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/194405/>

Version: Accepted Version

Proceedings Paper:

Singh, I., Bontcheva, K. orcid.org/0000-0001-6152-9600, Song, X. et al. (1 more author) (2022) Comparative analysis of engagement, themes, and causality of Ukraine-related debunks and disinformation. In: Hopfgartner, F., Jaidka, K., Mayr, P., Jose, J. and Breitsohl, J., (eds.) Social Informatics: 13th International Conference, SocInfo 2022, Glasgow, UK, October 19–21, 2022, Proceedings. 13th International Conference, SocInfo 2022, 19-21 Oct 2022, Glasgow, UK. Lecture Notes in Computer Science . Springer International Publishing , pp. 128-143. ISBN 9783031190964

https://doi.org/10.1007/978-3-031-19097-1_8

This version of the paper has been accepted for publication, after peer review (when applicable) and is subject to Springer Nature's AM terms of use, but is not the Version of Record and does not reflect post-acceptance improvements, or any corrections. The Version of Record is available online at: http://dx.doi.org/10.1007/978-3-031-19097-1_8

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Comparative Analysis of Engagement, Themes, and Causality of Ukraine-Related Debunks and Disinformation

Iknoor Singh^[0000-0002-3788-3295], Kalina Bontcheva^[0000-0001-6152-9600],
Xingyi Song^[0000-0002-4188-6974], and Carolina Scarton^[0000-0002-0103-4072]

Department of Computer Science, University of Sheffield, Sheffield (UK)
{i.singh, k.bontcheva, x.song, c.scarton}@sheffield.ac.uk

Abstract. This paper compares quantitatively the spread of Ukraine-related disinformation and its corresponding debunks, first by considering re-tweets, replies, and favourites, which demonstrate that despite platform efforts Ukraine-related disinformation is still spreading wider than its debunks. Next, bidirectional post-hoc analysis is carried out using Granger causality tests, impulse response analysis and forecast error variance decomposition, which demonstrate that the spread of debunks has a positive impact on reducing Ukraine-related disinformation eventually, albeit not instantly. Lastly, the paper investigates the dominant themes in Ukraine-related disinformation and their spatiotemporal distribution. With respect to debunks, we also establish that around 18% of fact-checks are debunking claims which have already been fact-checked in another language. The latter finding highlights an opportunity for better collaboration between fact-checkers, so they can benefit from and amplify each other's debunks through translation, citation, and early publication online.

Keywords: Disinformation · Debunks · Ukraine-related disinformation · Comparative analysis · Social media

1 Introduction

Following on from and interleaved with the COVID-19 infodemic, the war in Ukraine has unleashed a new large stream of mis- and disinformation [1], as evidenced, amongst others, by fact-checkers from the European Digital Media Observatory (EDMO) who found a record-high Ukraine-related disinformation in March 2022¹. Examples include viral decontextualised videos from past² and a popular pro-Kremlin false narrative about the existence of a biolab in Ukraine funded by Joe Biden's son³. To counter this fast-flowing disinformation, the International Fact-checking Network (IFCN) fact-checkers are working together

¹ <https://edmo.eu/fact-checking-briefs/>

² <https://www.politifact.com/factchecks/2022/may/10/facebook-posts/no-not-footage-ukraine-shooting-down-russian-plane/>

³ <https://www.politifact.com/article/2022/apr/01/facts-behind-russian-right-wing-narratives-claimin/>

to maintain and publish a unified database of debunks of Ukraine-related disinformation⁴. In order to measure the effectiveness of these efforts, we carry out a comparative analysis of engagement, themes, and predictive causality of Ukraine-related debunks and disinformation.

The novel contributions of this paper are in answering the following three key research questions through a quantitative analysis of Ukraine-related disinformation and debunks on Twitter:

- RQ1** What is the overall engagement of Ukraine-related disinformation and debunks on Twitter (Section 4)?
- RQ2** Does the spread of debunks have a positive impact in reducing Ukraine-related disinformation (Section 5)?
- RQ3** What are the underlying themes in Ukraine-related disinformation and their spatiotemporal characteristics on Twitter (Section 6)?

In the following sections, we will discuss first related work (Section 2) and then detail the data acquisition methodology for this study (Section 3).

2 Related Work

Ukraine-related pro-Kremlin disinformation [29] is not new [1, 14, 15]. For instance, Lange-Ionatamishvili et al. [14] and Mejias and Vokuev [15] studied the spread of disinformation on social media after the 2014 annexation of Crimea by the Russian Federation, while Erlich and Garner [7] investigated if Ukrainian citizens are able to discern between factual information and pro-Kremlin disinformation. Another study [8] investigated the effectiveness of Russian propaganda in swaying the views of its readers. Recently, Park et al. [18] released a Ukraine-related dataset of tweets and carried out an analysis of public reactions to tweets by state-affiliated and independent media. Miller et al. [16] studied the spread of tweets related to hashtags that were trending in February 2022. Nonetheless, these studies do not focus specifically on comparing Ukraine-related disinformation and debunks in terms of engagement, inter-relationship, and topics.

Prior literature on the spread of true and false information is extensive [10, 21, 28]. Nevertheless, this paper is related to prior work that studied the spread and dynamics of false information and debunks on Twitter [2–5, 13, 17, 19, 20, 23, 26, 30]. In particular, Burel et al. [5] compared COVID-related misinformation and fact-checks using impulse response modelling, causal analysis, and spread variance analysis, while Chen et al. [6] investigated the reasons why people share fact-checks and ways to encourage this further. Also, Siwakoti et al. [24] showed that user engagement with fact-checks increased significantly as a result of the COVID-19 pandemic. However, to the best of our knowledge, no study has examined the predictive causality between Ukraine-related disinformation and debunks, or their spatiotemporal characteristics and top disinformation themes.

⁴ <https://ukrainefacts.org/>

Table 1. Top domains of disinformation and debunk links.

Disinformation Domains	Debunk Domains
facebook.com (30%)	dpa-factchecking.com(25%)
tiktok.com (3%)	euvdisinfo.eu (25%)
twitter.com (3%)	rumorscanner.com (9%)
oroszhirek.hu (2%)	politifact.com (8%)
sputniknews.com (2%)	factly.in (8%)
nabd.com (2%)	verify-sy.com (6%)
arabic.rt.com (1%)	factcrescendo.com (4%)
fb.watch (1%)	verafiles.org (3%)
de.news-front.info (1%)	factcheck.org (2%)
Other (55%)	Other (10%)

3 Data

The data underpinning our analyses spans disinformation and debunks posted between 1 February and 30 April 2022. Specifically, we focus on Ukraine-related debunks and accompanying links to the corresponding disinformation encompassing: (i) 110 debunks and 311 links to disinformation published by EUvsDisinfo⁵, which primarily fact-checks pro-Kremlin disinformation; (ii) 344 debunks indexed by Google in the ClaimReview format⁶, which refer to 439 disinformation links. See Appendix A.1 for details on how we collect the disinformation links from debunks. Similar to Burel et al. [5], in addition to the above date restrictions, we also applied keyword-based filtering⁷ in order to select only Ukraine-related debunks and disinformation.

In total, this study analyses 454 debunk URLs and 750 links to Ukraine-related disinformation. The latter are provided by the fact-checking organisations themselves within the published debunks (see Appendix A.1), therefore we consider them as accurate. Table 1 shows the top domains that occur within the disinformation and debunk links. The former point either to content on social media platforms or to Kremlin-backed outlets. For debunks, the main domains are EUvsDisinfo (25%) and Dpa-factchecking (25%).

Next, we use academic research access to the Twitter API⁸ to obtain 16,549 unique tweets containing one of the above debunk URLs and another 62,882 unique tweets sharing one of the disinformation links⁹. Retweets are also collected, since we aim to investigate the overall spread of information on Twitter. Hereafter, the tweets containing debunk links are referred to as “debunk tweets”

⁵ <https://euvdisinfo.eu/>

⁶ <https://www.datacommons.org/factcheck/download>

⁷ Where debunked claims were in languages other than English, these were translated automatically with Google Translate first, prior to filtering with the keywords listed here: <https://gist.github.com/greenwoodma/430d9443920a589b6802070f2ca54134>

⁸ <https://developer.twitter.com/en/docs/twitter-api>

⁹ The dataset used for analysis received ethical approval from the University of Sheffield Ethics Board. This paper only discusses analysis and results in aggregate data, without providing examples or information about individual users.

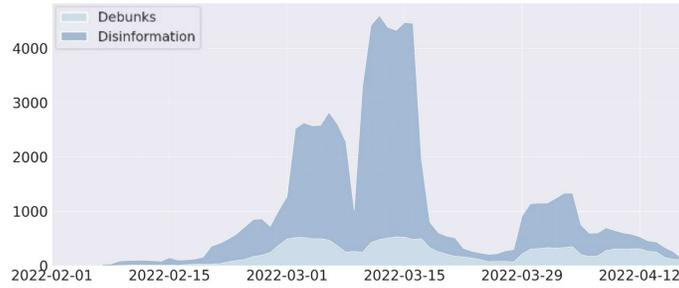


Fig. 1. Stacked plot of a rolling 7-day average of the number of disinformation and debunk tweets.

Table 2. Mean and standard deviation (STD) values of metrics of engagement with disinformation and debunk tweets. * represents a statistical significant difference ($p \leq 0.01$)

	Followers	Tweets	Retweets	Replies	Likes	Quote count
Mean - Disinformation	6,814	61,331	15.0*	0.5	4.5	0.2
Mean - Debunks	21,790*	89,098*	1.8	0.4	2.0	0.1
STD - Disinformation	2,04,422	1,22,527	312.3	34.5	631.9	13.9
STD - Debunks	4,55,884	1,55,183	15.0	7.6	28.0	1.3

and those containing disinformation links as “disinformation tweets”. Figure 1 shows the stacked plot of a rolling 7-day average curve for the spread of disinformation and debunk tweets. It shows that Ukraine-related disinformation spiked in the first half of March 2022, which consequently lead to an increase in published debunks as it is also reported in EDMO’s Fact-checking Briefs ¹⁰.

4 Comparative Analysis of Engagement

In order to measure the spread of disinformation and debunks through tweets, we first compare the differences in engagement metrics in terms of mean and standard deviation. Table 2 shows the statistics for author’s followers, author’s tweets, number of retweets, replies, likes and the quote count. We find that the number of retweets, replies, likes and the quote count are comparatively higher for disinformation. However, a t -test reveals statistically significant difference ($p \leq 0.01$) only for the number of retweets, i.e. significantly more Twitter users are retweeting posts containing disinformation URLs than debunk ones. There is also a statistically significant difference ($p \leq 0.01$) in the number of followers and tweet counts for users sharing debunks as opposed to disinformation. This is as expected since the former are primarily Twitter accounts of fact-checking organisations which naturally have more followers and post more frequently.

Figure 2 shows the histogram and kernel density estimate depicting the average number of days between the date of publication of disinformation tweets

¹⁰ <https://edmo.eu/fact-checking-briefs>

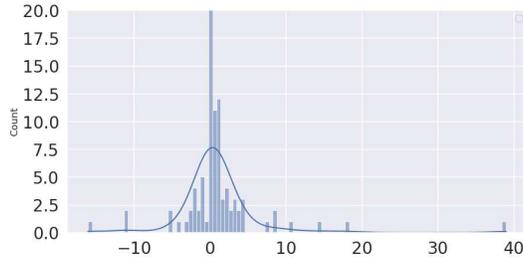


Fig. 2. Average difference in days between the date of publication of disinformation tweets and their corresponding debunk article (Fisher-Pearson coefficient of skewness = 3.37).

Table 3. Top ten cases with country affected by the disinformation and country of the authors of disinformation tweets.

Affected country	Authors' country	Percentage
Ukraine	Russia	9.0
Ukraine	Germany	7.0
Russia	Russia	6.0
Russia	Germany	5.0
Ukraine	United States	4.0
Ukraine	Venezuela	4.0
Ukraine	Mexico	3.0
United States	Mexico	3.0
Other		59.0

and their corresponding debunk articles. In this, for each debunk we compute $\sum_{i=1}^{|N|} (DoP_i - DoP_{debunk})/|N|$, where DoP_{debunk} is the date of publication of a debunk by the fact-checking organisation, DoP_i is the date of publication of a disinformation tweet i and $|N|$ is the total count of disinformation tweets for each debunk. The data is positively skewed, with a Fisher-Pearson coefficient of skewness of 3.37, suggesting some spread of disinformation even after the publication of the corresponding debunk article (see Section 5).

Since EUvsDisinfo debunks explicitly list countries where the disinformation is spreading, these can be compared to the country of the authors of those tweets. The latter is derived from the self-declared user location field obtained via the Twitter API¹¹ (when available). We obtain the location information for authors of 51% of the disinformation-sharing tweets. Unsurprisingly, the biggest proportion (9%) comes from cases where EUvsDisinfo has found the disinformation spreading in Ukraine, while the author's self-declared locations are in Russia (Table 3). Another key observation is the global nature of the disinformation, with spread extending significantly beyond Europe.

¹¹ <https://developer.twitter.com/en/docs/twitter-api/data-dictionary/object-model/user>. Where needed, Geopy Python library (Ref. <https://pypi.org/project/geopy/>) is used to extract the country name from the information provided by the API.

different languages. This finding demonstrates significant overlap in effort spent by fact-checking organisations in multiple countries, as well as cost- and time-saving opportunities that could be exploited with the help of translation and cross-publishing of debunks.

5 Post-hoc Causality Analysis

We test the bi-directional Granger causality [9] between the disinformation-sharing and debunk-sharing tweets. In other words, we want to investigate whether the spread of debunks has a positive impact on reducing the sharing of Ukraine-related disinformation on Twitter. Although identifying causation relationships between different information types is not trivial, a Granger causality test can be used to evaluate the predictive causality i.e. if the spread of one information type can be used to predict the spread of another. In this, we treat the occurrence of disinformation and debunk tweets as two time series variables and then try to find if one variable can be predicted from the other variable's past values and its own past values.¹⁴ First we build a Vector Regression model (VAR) [22], where a period of three is applied, based on the Akaike's Information Criterion. The Augmented-Dicky Fuller test identifies the data as stationary ($p \leq 0.01$). The general equation of VAR model is

$$disinfo(t) = \sum_{i=1}^k \alpha_{1,i} disinfo(t-i) + \sum_{i=1}^k \beta_{1,i} debunk(t-i) + \epsilon_1 \quad (1)$$

$$debunk(t) = \sum_{i=1}^k \alpha_{2,i} debunk(t-i) + \sum_{i=1}^k \beta_{2,i} disinfo(t-i) + \epsilon_2 \quad (2)$$

where $debunk(t)$ and $disinfo(t)$ refers to count of tweets at time t , k is the maximum lag order, $\alpha_{1,i}$ and $\alpha_{2,i}$ are autoregressive coefficients, $\beta_{1,i}$ and $\beta_{2,i}$ are regression coefficients, and ϵ_1 and ϵ_2 are error terms.

The experiments find a Granger causality relation, which shows that debunk spread has predictive causality over disinformation spread ($p \leq 0.01$). In addition, we also observe this weak causation in the opposite direction, i.e. from disinformation to debunks ($p \leq 0.01$). The significant results in both directions imply that changes in the spread of disinformation may induce changes in the spread of debunks and that the spread debunks may likewise cause changes in the disinformation spread. This is similar to the findings of previous work for COVID-19 misinformation [4, 5]. In order to further understand the weak causation between Ukraine-related disinformation and debunks, we use the VAR model to perform an impulse response analysis and forecast the error variance decomposition for 14 days periods.

Impulse response analysis is used to find the effect of shock in one variable to itself and the other variables in the VAR model. The prime reason to investigate

¹⁴ The Statsmodel Python library is used to perform the Granger causality test. Ref. <https://www.statsmodels.org/>

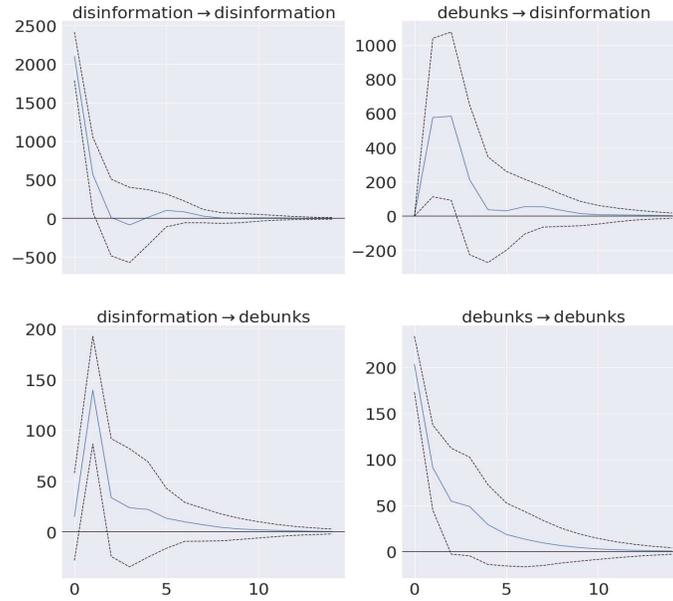


Fig. 5. Impulse Response Analysis (x-axis represents 14 days period and y-axis represents effect of shock). Top left and bottom left shows the effect of disinformation shock on disinformation and debunks respectively. Top right and bottom right shows the effect of debunk shock on disinformation and debunks respectively. By default, asymptotic standard errors are presented at the 95% confidence level.

this is to check if an increase in the spread of debunks triggers a reduction in disinformation on Twitter. Figure 5 shows that an orthogonal shock¹⁵ from debunks leads to an initial spike in disinformation but there is a downward trend for disinformation afterwards. This suggests that debunks will trigger a reduction in overall disinformation eventually, if not instantly. Similarly, an orthogonal shock from disinformation also triggers an initial spike in debunks (Figure 5) and an eventual decrease with time, although not instantly. This suggests swift response in debunk publication (mostly from fact-checkers) following a sudden rise in disinformation on social media. Interestingly, we also notice that the shock in disinformation quickly dies as the impact returns to zero with a sharp decrease on the second day, followed by a small post-shock peak during 4–6 day period which finally converges back to zero between 8–10 day period.

Forecast Error Variance Decomposition (FEVD) helps uncover the proportion of information each variable contributes in predicting a particular variable in the VAR model. FEVD analysis (Figure 6) reveals substantial predictive dependencies between Ukraine-related disinformation and debunks. Similar to what we find in impulse response analysis, FEVD results show that debunks directly affect disinformation by around 15% by the end of the 14-day period. We also

¹⁵ Cholesky decomposition is used for orthogonalisation

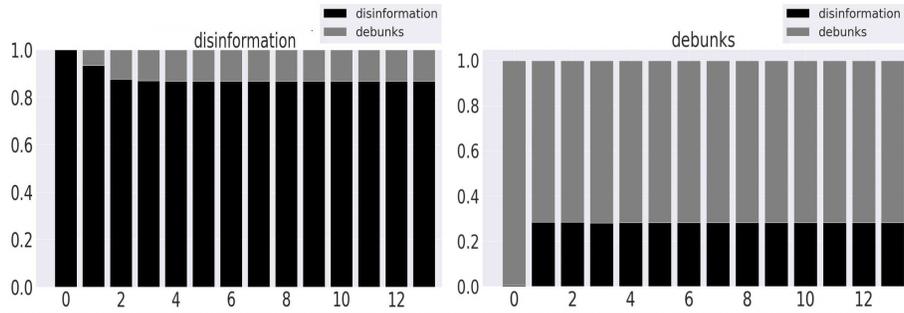


Fig. 6. Forecast Error Variance Decomposition (FEVD) plot (x-axis represents a 14-day period and y-axis represents proportion of affect). Left and right represents FEVD for disinformation and debunks respectively.

observe that debunks affect the spread of disinformation after an initial delay by a day, after which it rises and becomes constant following the third day. On the other hand, we also find that the spread of debunks is also affected by disinformation by around 30%. The results also show that the impact of disinformation on debunks is delayed initially for a day. In other words, this implies that the spread of debunks is not dependent on how Ukraine-related disinformation spreads initially. In summary, our experimental analysis confirms that the spread of debunking tweets does have a positive impact on reducing Ukraine-related disinformation on Twitter.

6 Topical Analysis of Ukraine-related Disinformation

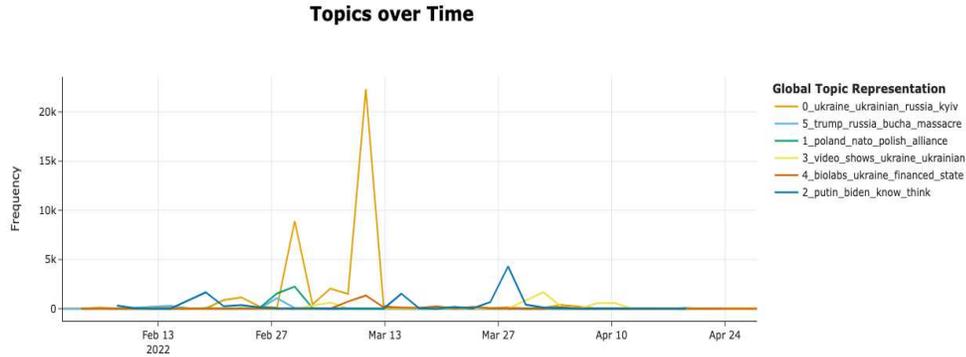
This section investigates the main topics in disinformation and study the engagement around them over time. The debunked claim statements are clustered by applying K-means to embeddings from the semantic search model¹⁶. The number of clusters is kept at six using the Elbow method and silhouette coefficient score. The model is run for a maximum of 300 iterations with K-means++ used as a method of initialisation. The clustering is applied on debunked claim statements and not on tweets itself. See Appendix A.1 for details on how we collect the debunked claim statements.

The class-based TF-IDF [11] is used to find top words in debunked claim statements in each of the clusters (Table 4). Each cluster has distinct words that separate it from the other five. In order to verify the separation between the clusters, we also plot a heatmap of topic similarity (see Appendix A.2). The results show that except clusters one and two, most of the clusters are distinct

¹⁶ We use the BERTTopic [11] Python library for clustering and MPNet [25] as the transformer model. Ref. <https://huggingface.co/sentence-transformers/all-mpnet-base-v2>

Table 4. Top ten words and count of disinformation tweets in each topic cluster. Order of word depicts its importance from left to right.

Topic clusters	Count
0.ukraine_ukrainian_russia_kyiv_neo_coup_nazis_war_crimea_weapons	39,761
1.poland_nato_polish_alliance_west_security_countries_western_europe_russian	10,225
2.putin_biden_know_think_vladimir_lee_answer_prices_says_oil	4,194
3.video_shows_ukraine_ukrainian_proof_jet_marcos_shot_soldiers_fighter	3,793
4.birolabs_ukraine_financed_state_biological_labs_military_biden_victoria_vaccinated	3,132
5.trump_russia_bucha_massacre_billions_evidence_west_100_planted_united	1,777

**Fig. 7.** Temporal spread of disinformation tweets in each topic cluster over time. Legend shows top five words of each cluster from Table 4.

in terms of the topics they cover. For instance, cluster four encompasses widespread conspiracies related to the U.S.-backed bioweapon labs in Ukraine and the US planning to send infected migratory birds to infect Russia¹⁷. Another cluster (one), includes the ongoing false narrative about the NATO country alliance being the real threat to Russia¹⁸. Table 4 also shows the count of corresponding disinformation tweets (and retweets) in each cluster, identifying that most of the tweets belong to cluster zero and one.

Next, we look at the temporal distribution of the disinformation tweets for each topic cluster. Figure 7 illustrates the line plot for topic prevalence between February and April 2022. For instance, cluster zero includes false claims related to Russia attacking Ukraine, Kyiv, neo-nazism, etc. and has two dominant peaks, one in the first week of March and another one in the second week of March (the tallest one at 10 March 2022). There is also an uptick in February suggesting that the disinformation narratives started spreading even before 24 February

¹⁷ <https://euvsdisinfo.eu/report/the-us-plans-to-send-infected-migratory-birds-to-infect-russia>

¹⁸ <https://www.politifact.com/factchecks/2022/feb/28/candace-owens/fact-checking-claims-nato-us-broke-agreement-again/>

and then spiked later in March. Similar results were found in the EDMO’s Fact-checking Briefs for February¹⁹ where they noticed a sudden increase in posts about growing tensions between Russia and Ukraine.

Cluster one (disinformation related to NATO and western countries) spiked in the first week of March. It includes a dominant false narrative about NATO attacking countries illegitimately²⁰.

Cluster two has multiple spikes: one in mid February; another one in mid of March; and the biggest one – at the end of March. It comprises false narratives involving Joe Biden, Vladimir Putin and Russian oil, e.g. that Biden’s cancellation of the Keystone pipeline “dramatically increased Americans” dependence on Russian oil.

Cluster three comprises of videos spreading disinformation and their distribution is fairly stable, with only a slight increase in the first and last week of March.

Cluster four contains conspiracies, such as an alleged presence of US biological labs in Ukraine²¹ and release of infected migratory birds to infect Russia²². Figure 7 shows that these type of conspiratorial narratives spiked during the second week of March. This is also coherent with the findings of EDMO’s Fact-checking Briefs for March 2022.

Cluster five contains disinformation related to Trump and the Bucha massacre and has comparatively time-limited span, with only a small peak at the end of February 2022.

7 Limitations and Future Work

Our work should be seen in the light of the following limitations. First, as described in Section 3, the study uses only tweets which contain explicit links to known disinformation or debunk articles. While this makes the dataset highly accurate and does not require additional human annotation, it also means that tweets that spread false claims or debunk them without citing a reference link could not be included. Second, this paper only discusses results on aggregate data, without looking at whether the tweets are from real or bot accounts. Lastly, user data, such as their country, is dependent on self-declared information in the user profiles, which is missing for many tweets. Nevertheless, the sample size is sufficiently large and robust to yield useful insights.

In future, we want to analyse the spread of disinformation and debunks before and after the start of the Russia-Ukraine war. We might have different answers for the research questions raised in the paper, which would potentially provide some insights into how an emergency event changes the spreading paradigms

¹⁹ <https://edmo.eu/fact-checking-briefs>

²⁰ <https://euvsdisinfo.eu/report/nato-is-not-a-defensive-alliance-it-attacks-countries-illegitimately>

²¹ <https://www.bbc.co.uk/news/60711705>

²² <https://euvsdisinfo.eu/report/the-us-plans-to-send-infected-migratory-birds-to-infect-russia>

of disinformation and debunks. We also want to find ways to automatically detect disinformation tweets which don't explicitly mention links or where the disinformation links mentioned are different from the ones present in our dataset. Lastly, the Granger causality test deals with linear relationship. Hence, in future, we plan to experiment with other non-linear tests like Hiemstra and Jones non-linear Granger causality [12] and Convergent Cross Mapping test [27].

8 Conclusion

This study carried out a comparative analysis of the spread of Ukraine-related false claims and debunks on Twitter between February and April 2022. In particular, our comparative engagement analysis found that tweets spreading disinformation are shared and retweeted significantly more as compared to those containing debunks. With respect to debunks, we also established that around 18% are focused on false claims for which debunks have already been posted in a different country or language. This finding is particularly important, as it points out to two opportunities going forward. Firstly, since many platforms, such as Facebook, already offer machine translation tools to their users, they could use that technology themselves to translate and match debunks automatically, so a false narrative spreading in one language can be flagged as false, based on an authoritative fact-check in another language. Secondly, fact-checkers themselves can benefit from using cross-lingual search and machine translation technologies to find such debunks, which they can then cite as a source or re-publish in translation and thus reduce the time elapsed between a false narrative starting to spread widely online and the time their debunk is published.

Another key finding is that the publication of debunks does ultimately lead to limiting the spread of Ukraine-related disinformation, albeit not immediately. In addition, FEVD results show substantial predictive dependencies between the spread disinformation and debunk tweets. Lastly, our data-driven analysis uncovered also the dominant themes in Ukraine-related disinformation and their temporal intensity. In conclusion, these findings have immediate relevance for a wide range of stakeholders, including digital platforms, fact-checkers, and online information users. The dataset used for analysis is available at <https://doi.org/10.5281/zenodo.6992686>.

Acknowledgements. This research has been partially supported by a European Union – Horizon 2020 Program, grant no. 825297 (WeVerify), the European Union – Horizon 2020 Program under the scheme “INFRAIA-01-2018-2019 – Integrating Activities for Advanced Communities” and Grant Agreement n.871042 (“SoBigData++: European Integrated Infrastructure for Social Mining and Big Data Analytics” (<http://www.sobigdata.eu>)).

References

1. Aguerri, J., Santisteban, M., Miró-Llinares, F.: The fight against disinformation and its consequences: Measuring the impact of “Russia state-affiliated media” on Twitter. *SocArXiv* (2022)
2. Allcott, H., Gentzkow, M.: Social media and fake news in the 2016 election. *Journal of economic perspectives* **31**(2), 211–36 (2017)
3. Barrera, O., Guriev, S., Henry, E., Zhuravskaya, E.: Facts, alternative facts, and fact checking in times of post-truth politics. *Journal of Public Economics* **182**, 104123 (2020)
4. Burel, G., Farrell, T., Alani, H.: Demographics and topics impact on the co-spread of COVID-19 misinformation and fact-checks on Twitter. *Information Processing & Management* **58**(6), 102732 (2021)
5. Burel, G., Farrell, T., Mensio, M., Khare, P., Alani, H.: Co-spread of misinformation and fact-checking content during the COVID-19 pandemic. In: *International Conference on Social Informatics*, pp. 28–42, Springer (2020)
6. Chen, Q., Zhang, Y., Evans, R., Min, C.: Why Do Citizens Share COVID-19 Fact-Checks Posted by Chinese Government Social Media Accounts? The Elaboration Likelihood Model. *International Journal of Environmental Research and Public Health* **18**(19), 10058 (2021)
7. Erlich, A., Garner, C.: Is pro-Kremlin Disinformation Effective? Evidence from Ukraine. *The International Journal of Press/Politics* p. 19401612211045221 (2021)
8. Gerber, T.P., Zavisca, J.: Does Russian propaganda work? *The Washington Quarterly* **39**(2), 79–98 (2016)
9. Granger, C.W.: Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: journal of the Econometric Society* pp. 424–438 (1969)
10. Grinberg, N., Joseph, K., Friedland, L., Swire-Thompson, B., Lazer, D.: Fake news on twitter during the 2016 us presidential election. *Science* **363**(6425), 374–378 (2019)
11. Grootendorst, M.: BERTopic: Neural topic modeling with a class-based TF-IDF procedure. *arXiv preprint arXiv:2203.05794* (2022)
12. Hiemstra, C., Jones, J.D.: Testing for linear and nonlinear granger causality in the stock price-volume relation. *The Journal of Finance* **49**(5), 1639–1664 (1994)
13. Jiang, M., Gao, Q., Zhuang, J.: Reciprocal spreading and debunking processes of online misinformation: A new rumor spreading–debunking model with a case study. *Physica A: Statistical Mechanics and its Applications* **565**, 125572 (2021)
14. Lange-Ionatamishvili, E., Svetoka, S., Geers, K.: Strategic communications and social media in the Russia Ukraine conflict. *Cyber war in perspective: Russian aggression against Ukraine* pp. 103–111 (2015)
15. Mejias, U.A., Vokuev, N.E.: Disinformation and the media: the case of Russia and Ukraine. *Media, culture & society* **39**(7), 1027–1042 (2017)

16. Miller, C., Inskip, C., Marsh, O., Arcostanzo, F., Weir, D.: Message-based Community Detection on Twitter (2022), URL <https://glavcom.ua/pub/pdf/49/4935/message-based-community-detection-on-twitter.pdf>
17. Nyhan, B., Reifler, J.: Estimating fact-checking's effects. Arlington, VA: American Press Institute (2015)
18. Park, C.Y., Mendelsohn, J., Field, A., Tsvetkov, Y.: VoynaSlov: A Data Set of Russian Social Media Activity during the 2022 Ukraine-Russia War. arXiv preprint arXiv:2205.12382 (2022)
19. Park, S., Park, J.Y., Chin, H., Kang, J.h., Cha, M.: An experimental study to understand user experience and perception bias occurred by fact-checking messages. In: Proceedings of the Web Conference 2021, pp. 2769–2780 (2021)
20. Recuero, R., Soares, F.B., Vinhas, O., Volcan, T., Hüttner, L.R.G., Silva, V.: Bolsonaro and the Far Right: How Disinformation About COVID-19 Circulates on Facebook in Brazil. *International Journal of Communication* **16**, 24 (2022)
21. Shao, C., Ciampaglia, G.L., Varol, O., Yang, K.C., Flammini, A., Menczer, F.: The spread of low-credibility content by social bots. *Nature communications* **9**(1), 1–9 (2018)
22. Sims, C.A.: Macroeconomics and reality. *Econometrica: journal of the Econometric Society* pp. 1–48 (1980)
23. Singh, I., Bontcheva, K., Scarton, C.: The false COVID-19 narratives that keep being debunked: A spatiotemporal analysis. arXiv preprint arXiv:2107.12303 (2021)
24. Siwakoti, S., Yadav, K., Bariletto, N., Zanotti, L., Erdogdu, U., Shapiro, J.N.: How covid drove the evolution of fact-checking. *Harvard Kennedy School Misinformation Review* (2021)
25. Song, K., Tan, X., Qin, T., Lu, J., Liu, T.Y.: Mpnet: Masked and permuted pre-training for language understanding. *Advances in Neural Information Processing Systems* pp. 16857–16867 (2020)
26. Swire, B., Berinsky, A.J., Lewandowsky, S., Ecker, U.K.: Processing political misinformation: comprehending the Trump phenomenon. *Royal Society open science* **4**(3), 160802 (2017)
27. Tsonis, A.A., Deyle, E.R., Ye, H., Sugihara, G.: Convergent cross mapping: theory and an example. *Advances in nonlinear geosciences* pp. 587–600 (2018)
28. Vosoughi, S., Roy, D., Aral, S.: The spread of true and false news online. *Science* **359**(6380), 1146–1151 (2018)
29. Yablokov, I.: Russian disinformation finds fertile ground in the West. *Nature Human Behaviour* pp. 1–2 (2022)
30. Zhang, Y., Guo, B., Ding, Y., Liu, J., Qiu, C., Liu, S., Yu, Z.: Investigation of the determinants for misinformation correction effectiveness on social media during COVID-19 pandemic. *Information Processing & Management* **59**(3), 102935 (2022)

A Appendix

A.1 Data Collection

As described in Section 3, we collect Ukraine-related debunks from EUvsDsinfo and ClaimReview. In order to collect the disinformation links, 1) the debunks indexed in ClaimReview schema has the `itemReviewed`²³ object which includes disinformation links that are being debunked by fact-checking organisation and debunked claim statement is present in `claimReviewed` object; 2) the debunks on EUvsDsinfo explicitly mention disinformation links on their website. Figure 8 shows the screenshot of one of the EUvsDsinfo debunks. The section enclosed in the red box contains disinformation links and the blue box represents the debunked claim statement.

The image shows a screenshot of a debunked claim statement and a list of disinformation links. The claim statement is in a blue box at the top, and the disinformation links are in a red box on the right. The main content is in a white box with a black header and an orange footer.

DISINFO: KYIV PREPARES CHEMICAL ATTACKS TO BLAME THE RUSSIAN ARMY

SUMMARY

Kyiv prepares for fabrications on Easter eve using chemical weapons in the city of Yuzhne to blame the Russian army for it

DISPROOF

Recurring pro-Kremlin disinformation narratives about false flag operations and chemical weapons attacks. Such claims are a part of a broader pro-Kremlin disinformation campaign accompanying the Russian invasion of Ukraine, trying to place the blame for the war's atrocities on Ukraine.

There is no evidence to back the claim that Ukraine is preparing a chemical or biological provocation or a use of the weapons of mass destruction.

PUBLICATION/MEDIA

- [almasryalyoum.com](#) (Archived)
- [mogaznews.com](#) (Archived)
- [arabyoum.news](#) (Archived)
- [arabic.rt.com](#) (Archived)
- [pressbee.news](#) (Archived)
- العالم الآن الدفاع الروسية: كيف تستعد لفكرات بأسلحة كيميائية في أوديسا عشية عيد الفصح أعيون الجزيرة الاخبارية (Archived)
- [almanar.com.lb](#) (Archived)
- [almaalomah.me](#) (Archived)
- [wiki-news.net](#) (Archived)

Fig.8. Screenshot of one of the EUvsDsinfo debunks. Section enclosed in the red box contains disinformation links and the blue box represents the debunked claim statement.

A.2 Heatmap

Figure 9 illustrates the heatmap of cluster similarity. The results show that except clusters one and two, most of the clusters are distinct in terms of the topics they cover. This indicates reasonable separation between the clusters found in Section 6.

²³ <https://schema.org/ClaimReview>

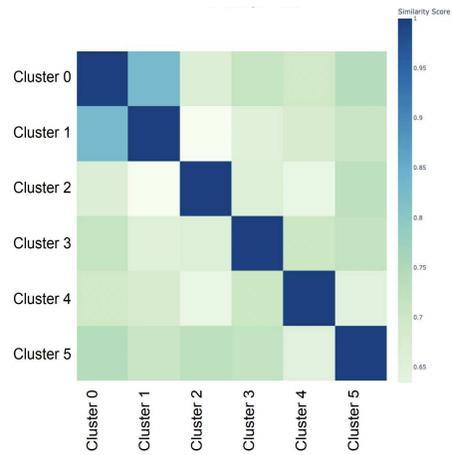


Fig. 9. Heatmap for topic cluster similarity. The description of clusters can be found in Section 6.