

Tweets in Time of Conflict: A Public Dataset Tracking the Twitter Discourse on the War between Ukraine and Russia

Emily Chen^{1,2}, Emilio Ferrara^{1,2,3}

¹ Department of Computer Science, University of Southern California

² Information Sciences Institute, University of Southern California

³ Annenberg School of Communication, University of Southern California
echen920@usc.edu, emiliofe@usc.edu

Abstract

On February 24, 2022, Russia invaded Ukraine. In the days that followed, reports kept flooding in from laymen to news anchors of a conflict quickly escalating into war. Russia faced immediate backlash and condemnation from the world at large. While the war continues to contribute to an ongoing humanitarian and refugee crisis in Ukraine, a second battlefield has emerged in the online space, both in the use of social media to garner support for both sides of the conflict and also in the context of information warfare. In this paper, we present a collection of nearly half a billion tweets, from February 22, 2022, through January 8, 2023, that we are publishing for the wider research community to use. This dataset can be found at <https://github.com/echen102/ukraine-russia>. Our preliminary analysis on a subset of our dataset already shows evidence of public engagement with Russian state-sponsored media and other domains that are known to push unreliable information towards the beginning of the war; the former saw a spike in activity on the day of the Russian invasion, while the other saw spikes in engagement within the first month of the war. Our hope is that this public dataset can help the research community to further understand the ever-evolving role that social media plays in information dissemination, influence campaigns, grassroots mobilization, and much more, during a time of conflict.

Introduction

Timeline of Key Events

The tensions between Ukraine and Russia have been on the rise for decades, but have substantially increased in the past few years. Following the Soviet Union's dissolution in 1991, Ukraine declared itself an independent country on August 24, 1991 (Sullivan 2022). In 2010, the pro-Russian presidential candidate, Viktor Yanukovich, was elected into office amid accusations of election fraud (Reuters 2022; Kahn 2022). Yanukovich was ousted in 2014 after backing out of an agreement with the European Union (EU), an agreement which would have brought Ukraine closer to becoming a member of the EU, in favor of engaging with Russia (Kahn 2022).

Shortly after Yanukovich was removed from office, Russia annexed Crimea from Ukraine against the wishes of

the West, further exacerbating tensions between Russia and Ukraine. Russia claimed that this action was supported by the vast majority of the Crimean people, but this was also denounced as a fraudulent vote by Western countries and Ukrainian leadership (Clinch 2022; Myers and Barry 2014; Collett-White and Popeski 2014).

Pro-Russian separatist groups in Donetsk and Luhansk in the Donbas region declared their independence, sparking the beginning of the war in Donbas (Reuters 2022). Several iterations of the Minsk agreement were signed between Ukraine and Russia in September 2014 in efforts to reach a resolution for the war; these efforts, however, have proven to be unsuccessful to this day (Sullivan 2022).

In 2019, the current President of Ukraine, Volodymyr Zelenskyy was elected to office, and one of his intentions was to end the conflict in Donbas (Pereira and Reeve 2022).

In late 2021 through early 2022, it became clear that Russia was amassing its forces near the Russian-Ukrainian borders. On February 21, 2022, Russia officially recognized the independence of Donetsk (Donetsk People's Republic) and Luhansk (Luhansk People's Republic) (Westfall 2022; Bloomberg 2022; Sullivan 2022).

On February 24, 2022, Russia invaded Ukraine, an act that drew swift condemnation from many world leaders, including the EU and NATO allies (McGee and Princewill 2022).

The Russian-Ukrainian war has caused and continues to be, as of these days, a humanitarian emergency for Ukrainians in the country, and has also created a refugee crisis as Ukrainians turned refugees flee to neighboring countries in large numbers (Schwartz, Kramer, and Gladstone 2022; Bloomberg 2022; Cengel 2022).

On March 2, 2022, Russia took control of Kherson, marking the first major Ukrainian city to fall to Russian troops (Schwartz and Pérez-Peña 2022).

As a result of the Russian invasion, many Western powers have imposed sanctions on Russia in an attempt to deter and reverse Russian aggression, and Western companies have begun to withdraw their operations from Russia (Furnakoshi, Lawson, and Deka 2022).

As of this writing (April 2023), the war has continued to rage on, with casualties on both sides of the conflict (Bigg 2023).

Social Media Activity

During this time of conflict, information warfare and campaigns continued, both in the lead-up and in the aftermath of the invasion. This has taken place on a variety of social media platforms, including Twitter. Russian disinformation campaigns have been rampant on social media (Badawy, Ferrara, and Lerman 2018; Broniatowski et al. 2018; Badawy et al. 2019; Badawy, Lerman, and Ferrara 2019; Dutt, Deb, and Ferrara 2019; Luceri, Giordano, and Ferrara 2020; Ferrara et al. 2020; Ezzeddine et al. 2023). Evidence suggests they are being carried out both domestically and abroad (Sharma et al. 2021; Scott 2022; La Gatta et al. 2023; Pierri et al. 2022). However, Ukrainians have also waged their social media fight against Russia, and Russian President Vladimir Putin, by using social media platforms to promote the Ukrainian cause and garner international attention and support towards their current plight (Cohen 2022; Garner 2022). Social media platforms have also taken actions to combat misinformation and disinformation in the wake of the conflict (Cohen 2022; Bushwick 2022; Ferrara 2022; Pierri, Luceri, and Ferrara 2022).

As the war continues to unfold, we hope that the research community can leverage our dataset to continue to combat misinformation, identify vulnerable communities and understand how the pervasiveness of social media has changed how modern conflict plays out both online and on the ground. These are just a few of many potential research directions that this data can help to answer, particularly in times when warfare is no longer limited by geographic bounds. In this paper, we document how to access this dataset and provide an overview of basic statistics and general findings from our Twitter data collection.

Data Collection

Our real-time Twitter data collection began on February 22, 2022. We used Twitter’s streaming API v1.1¹ to track keywords of interest that were both trending and related to the conflict at the time of collection. This version of their API has since been deprecated. This list of keywords primarily focused on terms related to events and locations prevalent in the conflict within the first several months of the war; however, we updated our list as needed. Our list of keywords is not exhaustive, but we did our best to monitor ongoing discourse and consulted knowledgeable colleagues. Prior work has shown that the streaming API is not completely random, which results in some biases depending on collection and collection volume (Mehrabi et al. 2021). However, we still leveraged the Twitter streaming API endpoint as it is the most sustainable method of gathering data that still gives us an understanding of current Twitter discourse.

We also leveraged Twitter’s search API to collect tweets prior to February 22, 2022; while we will continue to collect historical tweets, due to Twitter’s rate limit on the Academic Track search API, we are only able to use the search endpoint to collect 10 million tweets each month.² Our cur-

¹<https://developer.twitter.com/en/docs/twitter-api>

²<https://developer.twitter.com/en/products/twitter-api/academic-research>

rent data collection consists of over 570 million tweets, with more than 500 GB of raw data thus far. Our data collection concluded in late March 2023 due to recent modifications to the Twitter API. However, we remain optimistic that future API changes may enable us to recommence our data collection efforts in due course.

We host the dataset on our public GitHub repository so that other researchers are able to access tweet IDs that are pertinent to the Ukraine-Russia conflict, as misinformation and influence campaigns have already been detected in the narratives being pushed – both on and off Twitter. To remain in compliance with Twitter’s Terms & Conditions, we are not permitted to publicly release text or metadata pertaining to any specific tweet outside of their tweet IDs. As a result, we have published tweet IDs so that researchers can use the Twitter API or third-party tools, such as *Hydrator*³ or *Twarc*,⁴ to obtain the raw tweet payload. We note here that for any tweet that has been removed from Twitter (Pierri, Luceri, and Ferrara 2022), or if the account that posted a tweet has been deleted, suspended, or banned, researchers will be unable to retrieve the tweet content through Twitter’s API. Researchers can also ensure that their datasets are in compliance with Twitter’s Terms & Conditions through Twitter’s batch compliance endpoint to remove any tweets that have since been removed or hidden from public Twitter feeds.⁵

Tracked Keywords

We leveraged Twitter’s trending topics and hashtags to identify and curate a list of keywords that are pertinent to the developing war between Ukraine and Russia. We then used this list to query Twitter’s streaming API for any tweets that contained the keywords of interest in the tweet’s text. Twitter’s streaming API is capitalization agnostic, so we only need to track one capitalization permutation. We do, however, include some capitalization variations of keywords, particularly those that are not in English, to err on the side of caution. The list of the keywords that we tracked can be found in Table 1, but as events continued to develop and unfold, we expanded this list while consulting those more intimately familiar with the conflict. The most up-to-date keyword list can be found on our GitHub repository.

Data & Access Modalities

Release v1.2 (October 6, 2022)

Our third release contains tweets from the inception of our data collection, February 22, 2022, 4 AM UTC, through the end of October 1, 2022. This consists of a total of **454,488,445** tweets, and we detail general statistics about this release in this section. We note that we have since released v1.5 of this dataset, but for the purposes of this paper, we focus on Release v1.2. The language breakdown of release v1.2 can be found in Table 2, and the keywords that

³<https://github.com/DocNow/hydrator>

⁴<https://github.com/DocNow/twarc>

⁵<https://developer.twitter.com/en/docs/twitter-api/compliance/batch-compliance/quick-start>

Keywords	Tracked Since	Notes
ukraine	2/22/22	
russia	2/22/22	
putin	2/22/22	
soviet	2/22/22	
kremlin	2/22/22	
minsk	2/22/22	
ukrainian	2/22/22	
NATO	2/22/22	
luhansk	2/22/22	
donetsk	2/22/22	
kyiv	2/22/22	
kiev	2/22/22	
moscow	2/22/22	
zelensky	2/22/22	
fsb	2/22/22	
KGB	2/22/22	
Україна	2/22/22	Ukraine (uk)
Київ	2/22/22	Kiev (ru)
ФСБ	2/22/22	FSB (ru)
Россия	2/22/22	Russia (ru)
КГБ	2/22/22	KGB
Київ	3/1/22	Kiev (uk)
україни	3/1/22	ukraine (uk)
Росія	3/1/22	Russia (uk)
кгб	3/1/22	kgb (uk/ru)
фсб	3/1/22	fsb (ru)
SlavaUkraini	3/1/22	Glory to Ukraine (ru)
ukraini	3/1/22	ukraine
U+1F1FA, U+1F1E6	3/1/22	Ukrainian Flag emoji
Украина	3/1/22	Ukraine (ru)
украины	3/1/22	ukraine (ru)
Donbas	3/1/22	donbas
Донбас	3/1/22	donbas (uk)

Table 1: Keywords that we actively tracked (v1.2 — October 6, 2022). General notes and English translations provided in the notes section. Translations are followed with the original keyword language in parenthesis.

we were tracking are listed in Table 1. Some of the keywords were added later to our live collection which impacted the language proportions that we observed in our dataset; we used Twitter’s search API to obtain past tweets that we missed in our initial collection that also contain these keywords. We will be adding these additional tweet IDs in subsequent releases.

Language Distribution Twitter automatically attempts to tag each tweet with its language ISO and includes the found ISO in a tweet’s metadata. When we investigate the language distribution of the tweets, we find that English is the predominant language that is identified. This aligns with our expectations, as most of the keywords that we were initially tracking were all in English. When we added keywords in the Ukrainian and Russian languages on March 1, 2022, we saw a significant increase in the percentage of Ukrainian and Russian tweets in our dataset. We notice that the volume of tweets we collect initially averages over 4 million tweets collected each day, which we attribute to Twitter limiting the number of tweets we can collect due to high tweet volume (oso 2022); however, we see a gradual decline in the

Language	ISO	No. tweets	% total
English	en	321,088,619	70.65%
Spanish	es	18,358,931	4.04%
French	fr	17,857,397	3.93%
German	de	14,533,854	3.2%
Italian	it	11,589,565	2.55%
Undefined	(und)	11,473,234	2.52%
Russian	ru	9,968,421	2.19%
Japanese	ja	9,113,466	2.01%
Ukrainian	uk	8,016,384	1.76%
Turkish	tr	6,219,988	1.37%
Portuguese	pt	3,897,544	0.86%
Polish	pl	3,411,167	0.75%
Dutch	nl	1,837,698	0.4%
Indonesian	in	1,607,514	0.35%
Chinese	zh	1,430,735	0.31%

Table 2: The top 15 languages in our dataset and the number of respective tweets. (v1.2 — October 6, 2022)

volume of tweets we collect as the Ukraine-Russia war continued to rage on.

We also observe fluctuations in the volume of tweets posted in English, Ukrainian, and Russian, which, on an hourly granularity, matches the circadian patterns of the countries that predominantly speak that language and have a presence on Twitter. The English tweet volumes see an increase during general waking hours that follow United States time zones. This corroborates our finding, which we discuss later in this paper, that most tweets originate from the United States; Russian and Ukrainian tweets follow similar activity patterns, as Ukraine and Russia are in similar time zones.

Our data also shows that spikes of tweets in a particular language generally occur alongside major real-world events. When we examine the percentage of tweets that are written in Ukrainian and Russian, in general, Russian tweets make up a slightly larger percentage of the total number of tweets we collected on any given day. However, there are several instances where Ukrainian tweets outnumber Russian tweets. For example, on August 24, 2022, Ukrainians celebrated their independence day, which correlates with the increased percentage of Ukrainian tweets on that day (*cf.* Figure 1, bottom) (Hayda 2022; John and Kesaieva 2022).

Hashtags Table 3 lists the top 15 hashtags that we find used in our dataset. These hashtags are comprised of hashtags that are referring to groups involved or affiliated with the Ukraine-Russia war (e.g., #ukraine, #russia, #nato), and others are in support of Ukraine and against the war (e.g., #standwithukraine, #stoprussia). Other commonly used hashtags refer to locations in Ukraine that have been at the center of conflict during the war (e.g., #kyiv, #mariupol). The hashtag #tigray draws attention to the concurrent Tigray war and humanitarian crisis occurring in Ethiopia, with some wondering why the Ukrainian War has overshadowed the conflict in Ethiopia (Cheng and Anna 2022; Walsh 2022).

Outside of these hashtags, other highly prevalent hashtags include #anonymous, which referred to the announcement

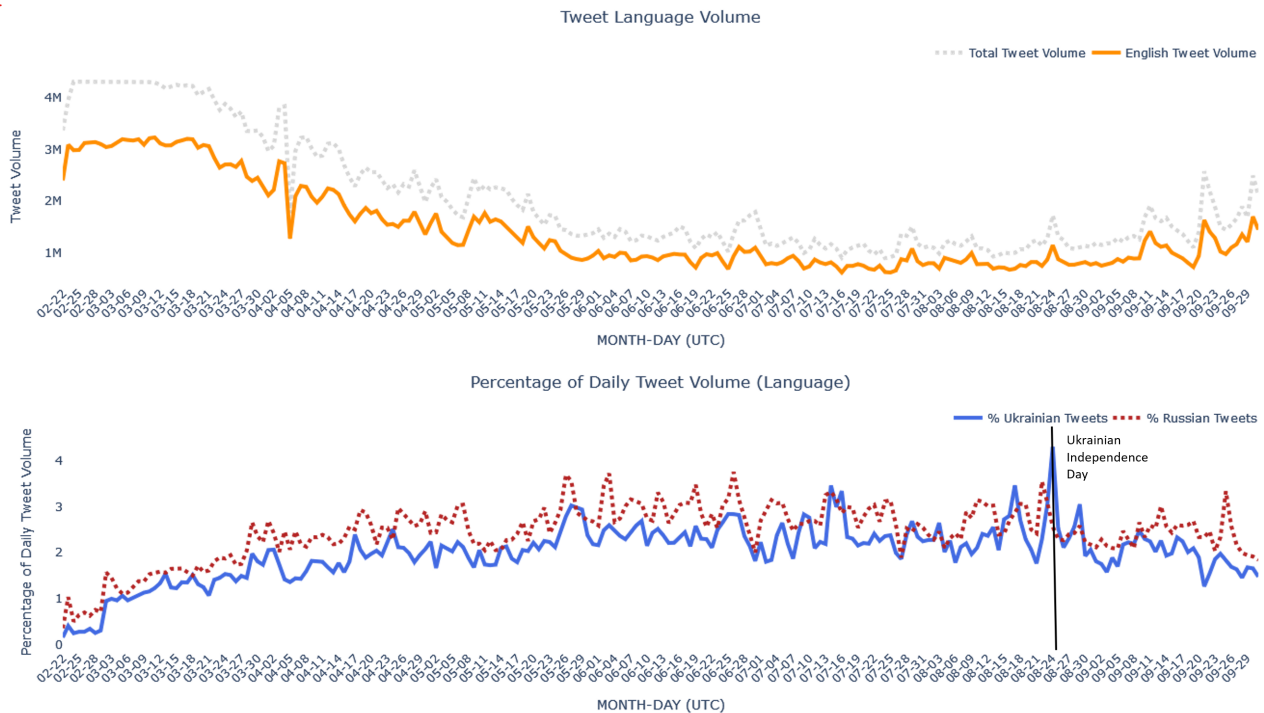


Figure 1: Total volume of all tweets collected and volume of English, Ukrainian, and Russian tweets in our dataset. Note that we observe a decline in overall tweet volume collected per day around March 12, 2022; tweet volume was limited by Twitter prior to this date due to the high volume of tweets using the keywords of interest. Engagement with these keywords began to wane in the following months.

Top 15 Hashtags

- ukraine
- russia
- putin
- standwithukraine
- ukrainerrussiawar
- nato
- russian
- ukrainian
- kyiv
- ukrainewar
- zelenskyy
- mariupol
- stoprussia
- slavaukraini
- tigray

Table 3: The top 15 hashtags used in our dataset, are listed in decreasing order from top to bottom. (v1.2 — October 6, 2022)

by the decentralized online hacking group Anonymous that they would be targeting Russian government websites and infrastructure while aiding Ukrainians (Pitrelli 2022).

When we compare the number of tweets using a hash-

tag containing the last name of the current political leaders of Russia and Ukraine, we can see that #putin is much more widely used than #zelenskyy. To further understand the magnitude of tweets that include #putin, we also plotted the number of tweets that use #biden for the current President of the United States, Joe Biden, and the number of tweets that use #trump, referring to the former President of the United States, Donald Trump (see Figure 2).

In Figure 2, we can see that #putin is used the most frequently, and experiences volatile usage which, in general, reacts to breaking news and events. There is a spike of engagement with Putin-related hashtags on March 4, 2022, and a manual inspection of these tweets reveal that many users were using the hashtag #SafeAirliftUkraine in tandem with #StopPutin, requesting the President of the United States, Joe Biden, to help with the evacuation of Ukrainians. This coincides with reports on March 4 that, despite discussions on temporary ceasefires the days before to establish humanitarian corridors to aid civilian evacuations, Russia was not complying with the agreements that were made (Stern et al. 2022; Stern, Suliman, and Taylor 2022). While hashtags using Biden and Trump see some use, the use of a hashtag mentioning Zelenskyy comes the closest to matching and sometimes slightly surpassing the number of tweets that use hashtags mentioning Putin. Given that Zelenskyy is the current president of Ukraine, these findings fall in line with our

Retweeted Country	→	Retweeter Country
United States	→	United States
United Kingdom	→	United Kingdom
Ukraine	→	United States
United Kingdom	→	United States
France	→	France
United States	→	Canada
United States	→	United Kingdom
India	→	India
Ukraine	→	Ukraine
Germany	→	Germany
Italy	→	Italy
Canada	→	United States
Turkey	→	Turkey
Ukraine	→	United Kingdom
Spain	→	Spain

Table 4: The top 15 locations for retweets where a location for both the retweeted and retweeter were identifiable. Locations are listed in decreasing order from top to bottom. (v1.2 — October 6, 2022)

Quoted Country	→	Quoter Country
United States	→	United States
United Kingdom	→	United Kingdom
United Kingdom	→	United States
France	→	France
Ukraine	→	United States
United States	→	United Kingdom
United States	→	Canada
Germany	→	Germany
India	→	India
United States	→	Japan
Japan	→	Japan
Spain	→	Spain
Russia	→	United States
Canada	→	Canada
Italy	→	Italy

Table 5: The top 15 locations for quoted tweets where a location for both the quoted user and user quoting a tweet were identifiable. Locations are listed in decreasing order from top to bottom. (v1.2 — October 6, 2022)

expectations.

Locations of Tweeters In each tweet’s metadata that Twitter’s API returns, we are given information about the user who posted a tweet. This will sometimes contain geolocation data from the user’s post, but in previous works, we have observed less than 1% of our dataset is returned with Twitter’s own geolocation information (Jiang et al. 2020). Thus, if a user manually provides a location in their user profile, we use this as a proxy to determine where a user is located (Jiang et al. 2020). While this is not a perfect method, this is one of the only ways we can ascertain a user’s general or affiliated physical location on a larger scale.

We find that most tweets originate from the United States

and the United Kingdom, which are primarily English-speaking countries. Given that the vast majority of our initial keywords were in English, this matches our expectations. There are two different ways a user can retweet a post – a user can simply retweet a post without adding additional commentary (we refer to this type of tweet as *retweets*) or a user can add additional text to the post (which we refer to as a *quoted* tweet). If user A retweets or quotes user B, we consider the direction of the retweet or quote to be from user B to user A. For quoted and retweeted tweets, we list the top 15 locations in Table 4 where we were able to identify a location for both the retweeted and the retweeter. We do the same for quoted tweets in Table 5.

Interestingly, we find that users in the United States quote users from Ukraine 3.34 times more than they quote users from Russia; users that are retweeting tweets from the United States are 9.11 times more likely to retweet users from Ukraine as opposed to users from Russia. Users from the United States retweeting users from Russia is the 25th most frequent location pairing.

Domain Sharing Finally, we take a look at the domains that users post as a part of their tweets. We consider any domain that is included within a user’s tweet as a domain that the user has shared, including domains that are part of a retweet. Within the top 50 domains that were shared, we find several domains that are known Russia state-sponsored media (www.RT.com), or tagged as questionable (www.rumble.com) and prone to conspiracies and pseudo-science (www.zerohedge.com) by Media Bias / Fact Check (MBFC) (Zandt 2022). MBFC is an independently operated website that classifies domains into certain factual and political lean categories, based on specific criteria that are described further on their per domain synopsis (Zandt 2022).

While the number of tweets that interact with these particular domains make up a small proportion of all tweets posted about the Ukraine-Russia war, we do see a fairly consistent number of tweets that interact with these domains over time (cf. Figure 3). ZeroHedge is the most consistently posted over our observation period, but we do note that these domains experience spikes in engagement occurred within the first month of the war, with Rumble being shared in the most tweets in early March.

Access

Our dataset can be found on GitHub at: <https://github.com/echen102/ukraine-russia>.

Due to recent changes to the Twitter API, we are no longer able to update this dataset with new data as of late March 2023. We hope to resume our collection and regular updates in the future if subsequent API changes allow us to do so. Given that monthly rate limits have significantly decreased, we suggest interested parties take a random sample of tweet IDs to hydrate over the desired time frame.

Researchers who would like to use our collected data must remain in compliance with Twitter’s Terms & Conditions,⁶

⁶<https://developer.twitter.com/en/developer-terms/agreement-and-policy>




Figure 2: Total volume of tweets that use the hashtags #biden, #trump, #putin and #zelensky. For #zelensky, we also counted #zelenskyy. The capitalization of a hashtag did not matter, as we lower-cased all hashtags.




Figure 3: Total volume of tweets that share specific domains in our dataset. RT.com is a known Russian state-owned media company, while MBFC has classified rumble.com as questionable and zerohedge.com as prone to conspiracy and pseudoscience.

and agree upon the terms of usage as outlined in the accompanying license.

Ethical Considerations

As mentioned earlier in this paper, we are only able to publish the tweet IDs associated with our collected tweets as a part of our data collection. This is to ensure that we remain in compliance with Twitter’s developer terms of service, which restrict us from publicly publishing any individual tweet information outside of the tweet’s unique ID. All tweets that can be retrieved by the end user are tweets that are publicly accessible – any tweet that has been deleted, made private, or was tweeted by a user that has since made their account private, was suspended, or deleted their account is no longer accessible.

The collection of this dataset was IRB-approved by USC.

Inquiries

If you have technical questions about the data collection, please refer to the issues section in the GitHub repository or contact Emily Chen at echen920@usc.edu.

For any further questions about this dataset, please contact Dr. Emilio Ferrara at emiliofe@usc.edu.

Acknowledgments

The authors gratefully acknowledge support from DARPA (contract no. HR001121C0169). The authors also thank Kristina Lerman, Jon May, and our lab for helping to identify pertinent keywords to track in our ongoing data collection.

References

2022. Suspicious Twitter Activity around the Russian Invasion of Ukraine. Technical report, Indiana University.

Badawy, A.; Addawood, A.; Lerman, K.; and Ferrara, E. 2019. Characterizing the 2016 Russian IRA influence campaign. *Social Network Analysis and Mining*, 9: 1–11.

Badawy, A.; Ferrara, E.; and Lerman, K. 2018. Analyzing the digital traces of political manipulation: The 2016 Russian interference Twitter campaign. In *2018 IEEE/ACM international conference on advances in social networks analysis and mining (ASONAM)*, 258–265. IEEE.

Badawy, A.; Lerman, K.; and Ferrara, E. 2019. Who falls for online political manipulation? In *Companion proceedings of the 2019 world wide web conference*, 162–168.

Bigg, M. M. 2023. Russia invaded Ukraine more than 10 months ago. here is one key development from every month

- of the war. <https://www.nytimes.com/article/ukraine-russia-war-timeline.html>. Accessed: 2023-01-15.
- Bloomberg. 2022. A Visual Guide to the Russian Invasion of Ukraine. <https://www.bloomberg.com/graphics/2022-ukraine-russia-us-nato-conflict/>. Accessed: 2023-01-15.
- Broniatowski, D. A.; Jamison, A. M.; Qi, S.; AlKulaib, L.; Chen, T.; Benton, A.; Quinn, S. C.; and Dredze, M. 2018. Weaponized health communication: Twitter bots and Russian trolls amplify the vaccine debate. *American journal of public health*, 108(10): 1378–1384.
- Bushwick, S. 2022. Russia’s information war is being waged on social media platforms. <https://www.scientificamerican.com/article/russia-is-having-less-success-at-spreading-social-media-disinformation/>. Accessed: 2023-01-15.
- Cengel, K. 2022. The 20th-century history behind Russia’s invasion of Ukraine. <https://www.smithsonianmag.com/history/the-20th-century-history-behind-russias-invasion-of-ukraine-180979672/>. Accessed: 2023-01-15.
- Cheng, M.; and Anna, C. 2022. WHO chief: Lack of help for Tigray Crisis due to Skin Color. <https://apnews.com/article/russia-ukraine-covid-health-09bce36f33719f793eaa79e1de520ad4>. Accessed: 2023-01-15.
- Clinch, M. 2022. How Russia invaded Ukraine in 2014. and how the markets tanked. <https://www.cnbc.com/2022/01/27/how-russia-invaded-ukraine-in-2014-and-how-the-markets-tanked.html>. Accessed: 2023-01-15.
- Cohen, R. 2022. A surge of unifying moral outrage over Russia’s war. <https://www.nytimes.com/2022/03/01/world/europe/zelensky-ukraine-war-outrage.html>. Accessed: 2023-01-15.
- Collett-White, M.; and Popeski, R. 2014. Crimeans vote over 90 percent to quit Ukraine for Russia. <https://www.reuters.com/article/us-ukraine-crisis/crimeans-vote-over-90-percent-to-quit-ukraine-for-russia-idUKBREA1Q1E820140316>. Accessed: 2023-01-15.
- Dutt, R.; Deb, A.; and Ferrara, E. 2019. “Senator, We Sell Ads”: Analysis of the 2016 Russian Facebook Ads Campaign. In *Advances in Data Science: Third International Conference on Intelligent Information Technologies, ICIT 2018, Chennai, India, December 11–14, 2018, Proceedings 3*, 151–168. Springer.
- Ezzeddine, F.; Luceri, L.; Ayoub, O.; Sbeity, I.; Nogara, G.; Ferrara, E.; and Giordano, S. 2023. Characterizing and Detecting State-Sponsored Troll Activity on Social Media. arXiv:2210.08786.
- Ferrara, E. 2022. Twitter Spam and False Accounts Prevalence, Detection and Characterization: A Survey. *First Monday*, (27).
- Ferrara, E.; Chang, H.; Chen, E.; Muric, G.; and Patel, J. 2020. Characterizing social media manipulation in the 2020 US presidential election. *First Monday*.
- Funakoshi, M.; Lawson, H.; and Deka, K. 2022. Tracking sanctions against Russia. <https://graphics.reuters.com/UKRAINE-CRISIS/SANCTIONS/byvrjenzmve/>. Accessed: 2023-01-15.
- Garner, I. 2022. How is the war going for Putin on social media? not great. <https://www.washingtonpost.com/outlook/2022/03/04/russian-social-media-ukraine-vk-propaganda/>. Accessed: 2023-01-15.
- Hayda, J. 2022. Kyiv hosts a different kind of parade to celebrate Ukraine’s Independence Day. <https://www.npr.org/2022/08/21/1118700800/ukraine-independence-day-russia-war>. Accessed: 2023-01-15.
- Jiang, J.; Chen, E.; Yan, S.; Lerman, K.; and Ferrara, E. 2020. Political Polarization Drives Online Conversations About COVID-19 in the United States. *HBET*, 2: 200–211.
- John, T.; and Kesaieva, Y. 2022. Ukraine’s Independence Day darkened by Deadly Missile Strike. <https://www.cnn.com/2022/08/23/europe/ukraine-independence-day-kyiv-display-intl/index.html>. Accessed: 2023-01-15.
- Kahn, J. 2022. Who is Viktor Yanukovich? the ousted Ukrainian president that Putin hopes to put back in power. {<https://fortune.com/2022/03/02/viktor-yanukovich-yanukovich-putin-put-back-in-power-ukraine-russia/>}, note = Accessed: 2023-01-15.
- La Gatta, V.; Wei, C.; Luceri, L.; Pierri, F.; and Ferrara, E. 2023. Retrieving false claims on Twitter during the Russia-Ukraine conflict. *WWW’23 Companion Proceedings*.
- Luceri, L.; Giordano, S.; and Ferrara, E. 2020. Detecting troll behavior via inverse reinforcement learning: A case study of Russian trolls in the 2016 us election. In *Proceedings of the international AAAI conference on web and social media*, volume 14, 417–427.
- McGee, L.; and Princewill, N. 2022. World leaders respond to Ukraine invasion, as fresh sanctions await Russia. <https://www.cnn.com/2022/02/24/europe/global-response-to-russia-ukraine-intl/index.html>. Accessed: 2023-01-15.
- Mehrabi, N.; Morstatter, F.; Saxena, N.; Lerman, K.; and Galstyan, A. 2021. A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)*, 54(6): 1–35.
- Myers, S. L.; and Barry, E. 2014. Putin reclaims Crimea for Russia and bitterly denounces the west. <https://www.nytimes.com/2014/03/19/world/europe/ukraine.html>. Accessed: 2023-01-15.
- Pereira, I.; and Reevell, P. 2022. What to know about Ukrainian President Volodymyr Zelenskyy. <https://abcnews.go.com/International/ukrainian-president-volodymyr-zelenskyy/story?id=83085078>. Accessed: 2023-01-15.
- Pierri, F.; Luceri, L.; and Ferrara, E. 2022. How does Twitter account moderation work? Dynamics of account creation and suspension during major geopolitical events. *arXiv preprint arXiv:2209.07614*.
- Pierri, F.; Luceri, L.; Jindal, N.; and Ferrara, E. 2022. Propaganda and Misinformation on Facebook and Twitter during

the Russian Invasion of Ukraine. *15th ACM Web Science Conference 2023*.

Pitrelli, M. B. 2022. Global Hacking Group Anonymous launches 'Cyber War' against Russia. <https://www.cnbc.com/2022/03/01/how-is-anonymous-attacking-russia-disabling-and-hacking-websites-.html>. Accessed: 2023-01-15.

Reuters. 2022. Timeline: The events leading up to Russia's invasion of Ukraine. <https://www.reuters.com/world/europe/events-leading-up-russias-invasion-ukraine-2022-02-28/>. Accessed: 2023-01-15.

Schwartz, M.; Kramer, A. E.; and Gladstone, R. 2022. Humanitarian crisis worsens for Ukrainians trapped in Russia's onslaught. <https://www.nytimes.com/2022/03/07/world/europe/ukraine-humanitarian-crisis-russia.html>. Accessed: 2023-01-15.

Schwartz, M.; and Pérez-Peña, R. 2022. First Ukraine City Falls as Russia strikes more civilian targets. <https://www.nytimes.com/2022/03/02/world/europe/kherson-ukraine-russia.html>. Accessed: 2023-01-15.

Scott, M. 2022. As war in Ukraine evolves, so do disinformation tactics. <https://www.politico.eu/article/ukraine-russia-disinformation-propaganda/>. Accessed: 2023-01-15.

Sharma, K.; Zhang, Y.; Ferrara, E.; and Liu, Y. 2021. Identifying Coordinated Accounts on Social Media through Hidden Influence and Group Behaviours. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 1441–1451.

Stern, D. L.; Firozi, P.; Paquette, D.; Pannett, R.; Francis, E.; Dixon, R.; Paúl, M. L.; Knowles, H.; and Kornfield, M. 2022. Limited cease-fire reached for civilian evacuations as Russian forces cut off Key Cities. <https://www.washingtonpost.com/world/2022/03/03/russia-ukraine-war-putin-news/>. Accessed: 2023-01-15.

Stern, D. L.; Suliman, A.; and Taylor, A. 2022. Russia not cooperating on proposed humanitarian corridor in Khereson, Ukrainian officials say. <https://www.washingtonpost.com/world/2022/03/04/kherson-ukraine-russia/>. Accessed: 2023-01-15.

Sullivan, B. 2022. Russia's at war with Ukraine. Here's how we got here. <https://www.npr.org/2022/02/12/1080205477/history-ukraine-russia>. Accessed: 2023-01-15.

Walsh, D. 2022. After secret U.S. talks fail, a hidden war in Africa rapidly escalates. <https://www.nytimes.com/2022/10/08/world/africa/ethiopia-tigray-war-talks-us.html>. Accessed: 2023-01-15.

Westfall, S. 2022. Why are Donetsk and Luhansk in Ukraine's Donbas region a flash point for Putin? <https://www.washingtonpost.com/world/2022/02/21/what-is-donbas-donetsk-luhansk-conflict/>. Accessed: 2023-01-15.

Zandt, D. M. V. 2022. Media Bias/Fact Check. <https://mediabiasfactcheck.com/>. Accessed: 2023-01-15.