

Russian Influence Operations during the Invasion of Ukraine

Joseph Littell and Nicolas Starck

Army Cyber Institute, West Point NY, United States

Joseph.Littell@WestPoint.edu

Nicolas.Starck@WestPoint.edu

Abstract: Prior to their invasion of Ukraine, the Russian Federation was seen as having a vastly superior ability to conduct operations in the information environment, particularly their ability to influence foreign audiences, when compared to their western counterparts (Cunningham, 2020). However, the efforts of Ukraine, NATO allies, and other aligned nations to conduct operations in the information environment with intelligence pre-bunking, traditional diplomacy, and sanctions, blunted many of Russia's best efforts. Further tactical and operational failures from the Armed Forces of the Russian Federation also undercut the salience of Russian messaging campaigns. This does not mean that Russia's efforts in the informational environment fully failed or did not adapt to these actions in the lead-up and continuation of the war. Through the use of the Natural Language Processing (NLP) technique transformer-based topic modeling (Grootendorst, 2022) and the causal inference technique Bayesian Structural Time Series analysis (Brodersen, 2015) this paper looks to both qualitatively and quantitatively examine how Russian state media on Twitter reacted, changed narratives, and focused efforts regionally from January 1st through September 1st, 2022. Through this analysis we argue that Russian efforts in Europe may have been of limited success. We further argue that by shifting focus, Russia gained influence in South America, and Middle East and North Africa, where their influence operations faced minimal obstructions, such as sanctions, and a latent anti-western sentiment.

Keywords: Influence Operations, Russia, Ukraine, Machine Learning, Natural Language Processing, Disinformation

Research Questions

1. How was Russia's state media messaging on Twitter oriented, both in frequency and content, prior to their invasion of Ukraine?
2. Was Russian messaging uniform, both in frequency and content, across languages and regions, or was it tailored by either?
3. Did Russian messaging change, either in frequency or content, in response to US and EU sanctions?
4. Did Russian messaging change, either in frequency or content, in response to changes in the physical/kinetic fight with Ukraine?

1. Background

The Russian Federation, like the Soviet Union preceding it, has used information in various ways to gain a strategic advantage both within its borders and with populations in foreign nations. Typically, this information is used to influence populations towards favorable positions for Russia. For example, during the First Chechen War Boris Yeltsin regularly lied to the Russian people. Starting as early as the day after Christmas, 1994, Yeltsin claimed that the military operations had ceased and Russia had successfully suppressed the Chechen rebellion (Hiatt, 1995). The historical record, however, proved that a further 18 months and tens of thousands of lives were lost before reaching that endpoint.

With improvements in telecommunications and the proliferation of the internet and social media, influence campaigns became much further reaching, with lower interaction costs and less impedance. Given their historical use of information, Russia was naturally one of the first nations to fully embrace the use of online influence campaigns as a mechanism to gain support for their military endeavors. This first materialized in a noticeable way during the annexation of Southern Ossetia from Georgia in 2008. Russia used a measured, now familiar, narrative of only acting defensively to protect ethnic Russians from Georgian aggression which was distributed through Russian outlets, to include web-based news and social media (Allison, 2008). The decentralized and pervasive nature of the campaign generated the sufficient ambiguity and plausible deniability that national governments and the international community turned a relatively blind eye to the incursion. The phenomenon of online information operations accelerated in 2014, again acting as a force multiplier to create the freedom of maneuver for Russia to conduct its invasion and annexation of Crimea.

Following the revelations that Russian influence campaigns targeted the 2016 US Presidential Election (National Intelligence Council, 2021) and the 2016 European Union Referendum in the United Kingdom (Bastos and Mercea, 2018), Russian online influence campaigns were seen as the pinnacle of information warfare. The disclosures relating to the 2016 US election led to several investigations which ultimately confirmed Russian

involvement. The investigations showed that Russia had directly targeted civilian populations, including influencing the conduct of local rallies by US citizens in support of their preferred political candidates. Regardless of the direct electoral efficacy of these influence operations on the US and UK populations, the operations created information effects. Numerous media outlets and think tanks opined that Russia was far outpacing their western counterparts' ability to conduct strategic information warfare, principally online influence campaigns. However, following the 2022 invasion of Ukraine much of the mystique and rhetoric surrounding Russian capability has faltered. Deliberate pre-bunking efforts and a unified NATO helped blunt Russia's efforts to sow doubt and dissuade Western support for the defense of Ukraine. In this paper, we present an analysis of the impacts of some of these counter-efforts and their effect on Russian information operations.

2. Methodology

The tweets used for this study were gathered by the Alliance for Securing Democracy for their Hamilton 2.0 Dashboard using the Twitter API (Schafer, 2019). A total of 611,860 tweets were compiled from 344 screen names associated with the Russian government. This set of screen names included accounts of news media sources such as Sputnik, RT, Tass, their various regional publications, and individual reporters associated with those outlets. The data sample includes 53 languages collected from January 1st, 2022 until September 1st, 2022. All tweets were machine translated to English from their native language through IBM Watson's translation.

We then filtered the data down to seven specifically selected languages: Russian, English, Spanish, French, Arabic, Turkish, and German. We chose these languages for four reasons. First, Russian, English, Arabic, Spanish, and Turkish, in order, are the five most prevalent languages within the data set. German and French round out the top ten languages, at nine and ten, respectively. Second, Arabic, Spanish, and to a lesser extent French, can be used to reflect Russian regional messaging in the Middle East, Central and South America, and Francophone Africa. Third, English, French, German, Spanish, and Turkish are all languages of NATO nations, a critical audience for the conflict. Finally, the Russian language allows for the examination of the messaging efforts about the invasion of Ukraine within the domestic population and the Russian diaspora. This filtered data set included 532,510 tweets from 341 accounts.

We used two approaches to answer the stated research questions. First to use quantitative methods to guide a qualitative assessment of the content, and second to quantify the frequency of the messaging. Both approaches used established machine learning techniques (Grootendorst, 2022) (Brodersen, 2015). As such, the methodology will summarize those methods and describe their application to answer our research questions.

We first conducted an analysis of the content of the filtered collection of tweets using Maarten Grootendorst's BERTopic: Neural topic modeling with a class-based TF-IDF procedure (Grootendorst, 2022). This methodology uses transformer based neural networks to convert sentences into a vector space that can be used to compare sentences semantically. Sentences are converted into an array of numbers that represent their position in some multi-dimensional space. This allows for the relationships between sentences to be measured by their distance to other distinct sentences within that same multi-dimensional space. Similar sentences will be closer in distance to each other, reflecting similar content. Further description of this process can be found in Devlin et al (2019) seminal paper introducing the Bidirectional Encoder Representations from Transformers (BERT) model. Comparing these large, multi-dimensional vectors can be computationally and temporally expensive. This can be mitigated by several dimensionality reduction tools within the field of machine learning. We selected Uniform Manifold Approximation and Projection for Dimension Reduction (UMAP) for its ability to retain global features projected onto a lower dimension. Once vectors are reduced, they are clustered based on the density in which they appear in the reduced space through a Hierarchical Based Density Scan (HBDSCAN) (Campello, 2013). This process allows for the identification of the most significant clusters within vector space. Finally, each density-based cluster is expressed as a human-understandable topic using Term Frequency Inverse Document Frequency (TF-IDF) to assign the most distinct words with a cluster as its identifier.

Our second approach applied a causal inference analysis to the frequency of Russian messaging on Twitter using Bayesian structural time-series models (Brodersen, 2015). We counted tweets by day to visualize the frequency of messaging over time. This daily count was measured from January 1st through September 1st, 2022. We then processed these daily counts through a diffusion-regression state-space model to calculate the expected quantity absent an intervention. This baseline projection allows for the inference of the temporal changes that are explained by the impact of a given intervention, the incorporation of Bayesian priors during the treatment, and to account for many sources of potential variation including the time-varying covariates. We can then

calculate the difference between the expected volume of Russian messaging given no intervention and the actual value to determine if there is statistically significant difference that can be correlated with the intervention. This metric provides a quantitative measure of the impact of these intervention strategies on Russian messaging. We then identified a selection of dates associated with interventions, such as waves of sanctions, changes to military strategy, and debt default, for which we could quantitatively measure the potential impact to Russian messaging. These events were selected based on the hypothesis that they would change the frequency of Russian activity on Twitter. While proving true causal relationships is not possible with the available data, this approach allowed us to quantify the effect of particular counter-influence actions relating to the Russian invasion of Ukraine.

3. Results

3.1 RQ1. How was Russia's state media messaging on Twitter oriented, both in frequency and content, prior to their invasion of Ukraine?

Prior to Russian Federation Forces crossing into Western Ukraine across multiple avenues of advance, Russian outlets and political leaders attempted to downplay the belief that Russia was amassing forces for a full-scale invasion. Russian Forces were being moved to the Ukrainian border throughout 2021, to such an alarming degree that on March 31st, 2021, the United States European Command raised their awareness level to indicate a potential crisis (Funaiole et al, 2021). By December 2021, Russia demanded a reduction of NATO presence in Central and Eastern Europe and a permanent restriction against Ukraine ever joining NATO (Wintour, 2022). These actions occurred immediately the start of our data sample on January 1st, 2022.

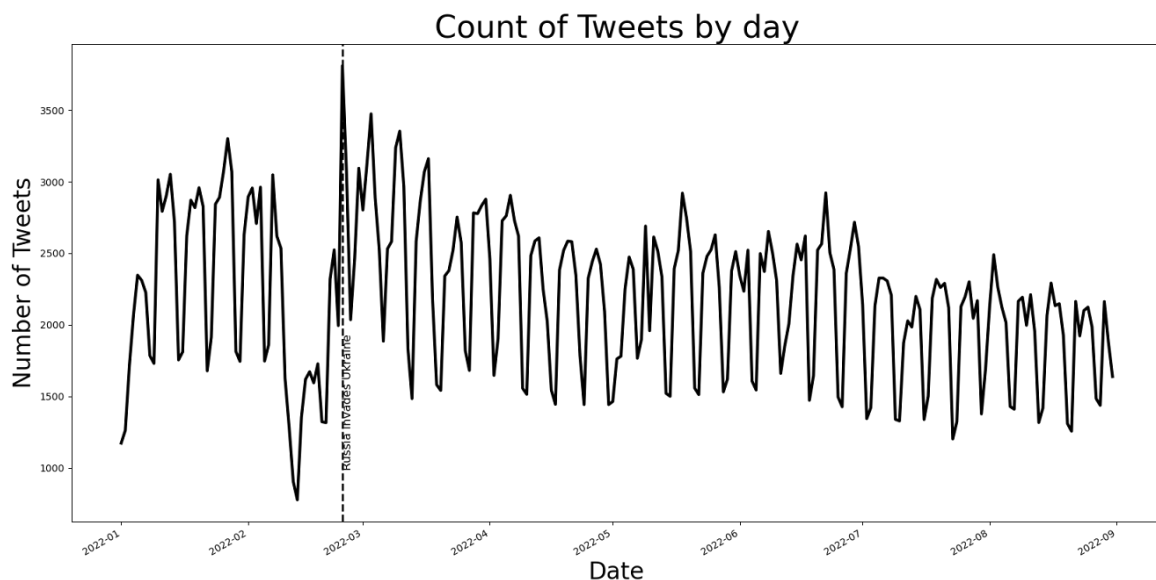


Figure 1: Frequency of Russian Media Tweets from January 1st through September 1st, 2022

The frequency of posts among all languages was stable until the February 4th, 2022, when all languages saw a sharp decline. The lowest volume of Russian Twitter activity was reached on February 13th, 2022, with only 778 tweets across all languages as compared to a daily average of 2,191 tweets. This decline occurred simultaneously with a series of key pronouncements from the West: (1) the deployment of 3,000 more NATO troops within Romania, Poland, and Germany, (2) coordinated public announcements of unified efforts to deter Russian aggression, and (3) the intent to impose severe economic sanctions if Russia intervened militarily in Ukraine. This decline, however, was not uniform across all languages. German language tweets increased during this period, and the decline in English language tweets was smaller than the other languages. This low point in Russian Twitter activity was followed by a sharp increase across all languages leading up to the invasion, with the highest volume, 3,810 tweets, occurring on February 24th. This represented a 74% increase on the average number of tweets across the dataset.

As with the frequency, there was a drastic shift in narrative and language in this same period preceding the invasion. Prior to the invasion there was no discernible consistent Russian narrative in any of the languages in

our sample. This pattern shifted around the launch of the “special military operation”. We now present a qualitative summary of NLP analysis by region, however the higher order trend that appeared was Russia trying to hide their intentions to invade Ukraine until approximately February 13th, at which point numerous regionally aligned narratives surfaced. These emergent narratives served one goal - to justify Russian military actions within Ukraine.

3.1.1 Russian

Within the Russian language tweets in the dataset preceding the invasion, we saw a range of regular news topics. This spanned domestic news such as COVID-19, the success of Russian-born athletes in the 2022 Winter Olympic Games, and international reporting on events like the Kazakhstan protests in early January and the Tonga Volcano and subsequent Tsunami. Notably absent was news or narratives relating to the massive troop buildup along the Ukrainian border.

However, in the surge of activity in the lead up to the invasion, Ukraine appeared as the central topic framed in two ways. The first framed the action as a large deterrence exercise aimed at showing the West Russia’s military might. The second framed a quick military intervention against Ukraine as critically needed due to their alignment with Nazis, and the needed recognition of the eastern Ukrainian regions of Lugansk and Donetsk as independent states.

3.1.2 Americas

Within the English-speaking world, the narrative predominantly focused on various COVID-19 topics including the Sputnik vaccine and its superiority to the western mRNA vaccines, Western vaccine mandates, and other anti-vaccine conspiracies. These COVID-19 narratives were coupled with news coverage of Julian Assange’s extradition, flooding in Brazil, and sporting events like the UFC. Mentions of Ukraine were sparse, only spiking once following US Secretary of State Anthony Blinken’s reaffirmation of Ukraine’s sovereignty being essential to global order (Zhang, 2022).

Following the decline in mid-February relative to the previous 6 weeks, the narrative changed drastically in the lead up to the invasion. Multiple major conspiracy theories appeared with the goal to add legitimacy to the forthcoming invasion or at worst, erode western support to Ukraine. The first stipulated that Dominion Voting Systems servers – allegedly used to rig the 2020 US Presidential election in favor of now-President Joseph R. Biden – were being held in Ukraine. The second conspiracy was that bioweapons, to include SARS-CoV-2, were engineered in US-backed labs across Ukraine. Another major conspiracy was that Ukrainian President Vladamire Zelenskyy was a puppet of the West, particularly the United States, and therefore an illegitimate leader. Finally, the persistent narrative that Ukraine has always been a part of Russia, the so-called country and culture is not real, and therefore Russia has historical rights to the land.

In Central and South America, the Spanish language tweets were again tailored to regional news. COVID-19 once again was a major source of content, as was traditional news including the revolt in Kazakhstan, an oil spill in Peru, and ongoing political strife and violence in Honduras. Like the English language tweets, US Secretary of State Anthony Blinken’s affirmation of Ukrainian sovereignty was refuted by the Russian media. However, the Spanish language messaging directly labeled the statement as disinformation and included a fact sheet to counter Secretary Blinken’s statement. The other mentions of Ukraine framed the Russian activity in several ways. First as a simple military exercise which should not warrant a NATO response, and second as necessary to evacuate Donetsk and Lugansk due to Ukrainian aggression. This pre-invasion rhetoric culminated with the announcement of a Special Military Operation to ensure the independence of the newly formed breakaway republics.

3.1.3 Middle East and North Africa

Arabic themes in January and early February centered around news and local events of the Middle East and North Africa (MENA). Topics included the ongoing conflict between Saudi Arabia and Houthis rebels in Yemen, the African Cup soccer tournament and Burkina Faso’s championship victory, ISIS reconstituting in Syria, and a looming Libyan government collapse. News of COVID-19, prevalent in many other languages’ tweets, were minimal within the Arabic language data. The Arabic themes mirrored the pattern seen in the Russian language tweets. Ukraine content was nearly non-existent prior to the decline in overall volume, followed by the narratives that Russia was a benign actor resisting Western aggression and the necessity for Russia’s recognition of Lugansk and Donetsk as independent states.

As Turkey sits at the crossroads of Asia and Europe, its news content was a mix of the MENA content and the COVID-19 narratives in the western language data. Another common narrative focused on Turkey’s dire

economic situation with rising unemployment, runaway inflation, and rising energy costs. No mention of Ukraine existed before the lead-up to the Russian Invasion, to include mentions of US Secretary of State Blinken's statements in January 2022. Curiously, mentions of Ukraine were still minimal in the lead up to the invasion. However, the messaging that was present followed the previously established narrative of protecting the newly independent Donbass republics from Ukrainian aggression.

3.1.4 Europe

French language tweets focused on traditional news, such as COVID-19 and travel-related content, with minimal mention of the Ukrainian or the troop build-up. The messaging centered on the then-upcoming French Presidential election prior to the invasion, potentially in the hopes of casting support for Ukraine as politically toxic within the French population to dissuade further support by political leaders. This lack of content relating to Ukraine remained consistent throughout the ramp up to the invasion.

The German language tweets ran contrary to all the other observations. The German tweets saw no decline in the frequency of posts in the run up to the invasion. The main narrative structure focused on the rising costs of energy and Nord Stream 2 construction. As with the assumption that political motivation may have deterred French actions, the economic costs and Germany's dependence on imported petroleum seemed to be the focal point for Russian messaging targeted at German audiences.

3.2 RQ2. Was Russian messaging uniform, both in frequency and content, across languages and regions, or was it tailored by either?

As demonstrated in the analysis of our first research question, Russia tailored its messaging towards specific regions. While Russia content was consistent in seeking ambiguity with regards to its invasion of Ukraine, the motivations and messaging justifying the necessity of the military intervention differed greatly by region. Regions outside of the West had their stories biased toward Western military aggression and imperialism. This framing had the potential to resonate in regions where distrust of the West is high or regions accustomed to Western military interventions, such as the Spanish, Arabic and Turkish speaking world. Russian speakers, while still susceptible to similar narratives, may have required further reasoning why military intervention, rather than standard deterrence, was needed. This could explain the rhetorical focus on Ukraine's supposed Nazism, the perpetration of Ukrainian war crimes in the Donbass War, and the importance of intervening militarily to ensure the safety and self-determination of the Russian-speaking regions of Ukraine.

This contrasts starkly with the Western messaging, which had vastly different narratives and rationales concerning the purpose of the invasion. The English language tweets focused on vaccine hesitancy, conspiracy, and US-specific political motifs. The German language messaging focused on industry and the economy, particularly the energy sector. Finally, the French content focused on appeals to freedom and independence from tyrannical Ukrainian rule. While the Western tweets paint NATO and the EU in imperialistic terms, it does not belabor the point. This is an intuitive observation, as the West would be unlikely to view their own history of interventionism as a negative. Therefore, such messaging would be counterproductive in the pursuit of dissuading a response in the defense of Ukraine.

3.3 RQ3. Did Russian messaging change, either in frequency or content, in response to US and EU sanctions?

The day following the February 21st, 2022, national television address by Russian President Vladimir Putin, Western nations announced their first round of economic sanctions against Russian political leaders and oligarchs. By the 26th of February, these sanctions escalated from individuals to Russian institutions including the banking sector, access to western social media, internet services, oil imports and exports, and more. While additional sanctions continued to be imposed throughout the conflict, the vast majority were in place within the first month of the conflict. The most temporally direct effect to Russian information campaigns came on February 27th, 2022, when the European Union member states banned all Russian state-owned outlets, such as RT and Sputnik (Siebold and Blenkinsop, 2022).

Taking the first five weeks of data collection as a baseline, and accounting for the lull in tweets from February 8th to the 23rd, Russia media overall saw a steady decline since the invasion in its output of content across all languages combined as shown in Figure 2 below. However, not every language decreased at the same rate nor degree. For example, English, French, and German all saw a significant, steady decline, while Arabic tweets had

a minimal decline, and Turkish and Spanish language tweets both had periods of increasing frequency following the invasion and requisite sanctions.




Figure 2: Impact of sanctions of Russian Media Tweets across all languages

During the post-intervention period, the average number of daily tweets was approximately 2,520. In the absence of any sanctions, we would have expected an average response of 2,740. In relative terms, the actual average number of daily tweets, our response variable, showed a decrease of 8% with the 95% interval of this percentage being [-17%, +1%]. This is a statistically significant finding as the likelihood of this happening by chance is very small, with the Bayesian one-sided tail-area probability $p = 0.048$. Comparing the cumulative effect against the French and German data, where these sanctions took place, the decline becomes even more clear. During the post-intervention period, French language tweets had a daily average value of approximately 34.20. In the absence of an intervention, we would have expected an average of 69.38 tweets per day. This is a 51% decline, with a 95% interval of this percentage is [-61%, -41%] and a Bayesian one-sided tail-area probability $p = 0.001$. German language tweets saw an even greater effect than French, with a 54% [-76%, -34%] decline with a Bayesian one-sided tail-area probability $p = 0.001$.

By contrast, for Spanish language tweets we would have expected an average response of 374.31 in the absence of an intervention. However, during the post-intervention period the response variable had an average daily value of approximately 352.91 daily tweets. This is only 6% lower than the prior expectation with a 95% confidence interval of [-17%, +5%]. Observing Figure 3 below, we can be seen that a drop off did not occur until July, potentially as a result of the official recognition on July 27th of Russia's debt default from the prior month (King and Jordan, 2022). Unlike the overall, French, and German examples, this finding is not statistically significant as it has a Bayesian one-sided tail-area probability $p = 0.148$. This is particularly notable compared to the Arabic dataset, another language unaffected by direct sanctions. The Russian activity on Arabic language Twitter saw only a 7% decline overall, with a 95% confidence interval of [-13%, -0%]. These findings, unlike those for Spanish, were statistically significant with a Bayesian one-sided tail-area probability $p = 0.019$.




Figure 3: Impact of sanctions of Russian Media on Spanish language tweets

Following the EU sanctions of Russian state-run media, two narratives appeared in the remaining French and German language tweets. The first narrative was a complaint about censorship of Russian media. This narrative pointed out the hypocrisy of the West in removing Russian media outlets access to these platforms and made claims that Western governments control Western social media platforms. The second narrative tried to migrate users from Twitter onto Telegram. This migration would allow the uncensored expression of Russian narratives, and interaction within the Telegram channel could be monitored and moderated by Russian actors. By contrast, many of the same narratives used in the lead up to the invasion persisted well into June 2022 within the English language data. The narrative around biological labs was particularly active throughout the spring and summer before being surpassed by a narrative blaming Ukrainian support for rising inflation.

3.4 RQ4. Did Russian messaging change, either in frequency or content, in response to changes in the physical/kinetic fight with Ukraine?

From the standpoint of frequency, we found no statistically significance date from a set of the following events when looking at the cumulative data set – the Bucha massacre (2022-04-01), the end of Kiev Offensive (2022-04-04), the launching of Eastern Offensive (2022-04-18), Finland and Sweden signal interest in joining NATO (2022-04-10), Finland and Sweden officially apply for NATO (2022-05-18), the Russian capture of Mariupol (2022-05-20), Russia failing to pay foreign debt (2022-06-24), or Russia officially declared to have defaulted on its debt (2022-06-27).

However, when separated by language, some trends do appear. As mentioned in RQ3, Spanish language saw a decline in posts following the Russian foreign debt default, which began on June 24th, 2022. This was a relative decline, visualized in Figure 4 below, of 16% with a 95% confidence interval of [-24%, -8%], with a Bayesian one-sided tail-area probability $p = 0.001$. Russian language activity saw a similar decline of 6% [-19%, +6%] although it was not statistically significant ($p = 0.176$). No other event causes a noticeable change nor one that would be statistically significant.




Figure 4: Impact of Debt Default on Spanish language tweets

For content, Russia references all of these events and used them as part of regionally based narratives. Within the English dataset, the messaging responded to the Bucha massacre claiming it as a hoax with actors and staged images. Within the Russian language data, the massacre was also labeled as a hoax. Notably this narrative also claimed if the event were not a hoax, a Nazi presence in Bucha were the real perpetrators. The messaging response to the capture of Mariupol also reflected subtle differences between languages. In Russian, the narrative focused on solidifying support for the “Special Military Operation,” fighting off the Nazis within the Azov Battalion, and ties to the WW2 victory against Nazi Germany. In English, however, the narrative focused on Ukrainians committing war crimes against the civilian population.

4. Conclusion

Numerous factors affected the Russian information campaign surrounding the invasion of Ukraine and its later narrative alignment towards distinct regional audiences. Previously, Russia had a great deal of informational freedom and were relatively uncontested when dealing with the West. In certain cases, such as the 2016 US Presidential Election and the 2016 European Union Referendum in the United Kingdom, Russia tailored its narratives to issues with deep seeded cultural significance, allowing it to compound existing social and cultural rifts. Russia’s information operations supporting military actions in Georgia and the 2014 annexation of Crimea existed in a vacuum, with limited push back from opposing informational forces. Any Western push back that did occur was limited and often not coordinated among allies, leaving large gaps for Russian narratives to permeate. Notably, the speed with which these military actions were concluded also limited the time for a response to develop.

However, Russia’s invasion of Ukraine faced a unified front from the West. This primarily occurred in the form of pre-bunking Russian narratives and levying heavy sanctions which hindered some informational capabilities as shown in RQ3. Many traditional military functions including logistics were also degraded. This forced Russia to rely on long shuttered factories, mothballed equipment, outside industry, and coordination with the few allies it had remaining. Some of the sanctions directly hindered capabilities by attacking Russian influence at its source within a given country, such as European Union member nations shutting down RT and Sputnik (King and Jordan, 2022). Other sanctions degraded non-informational aspects of the invasion, which we speculate may have caused Russia to shift support or funding away from informational efforts.

These sanctions packages, along with Russia’s own failures on the battlefield, slowed the pace of the war and eroded any previous first mover advantage Russia may have had. These delays also allowed for Ukraine to deploy a strong information capability to counter Russian narratives. Ukraine’s information campaign was bolstered by examples of both military victories and Russian war crimes. The long buildup of Russian forces and length of the

war also allowed traditional new organizations across the globe to send journalists and war correspondents to Ukraine. This effort added more visibility and first-hand content from the war that were not present during the quick invasions of Georgia in 2008 and Crimea in 2014. Taken as a whole, the global response to the Russian invasion of Ukraine in 2022 demonstrated an increased understanding and sophistication of the information environment. This should prompt a re-evaluation of the established view that Russia is the preeminent force in the information environment and serve as a roadmap for future effective responses to their tactics.

References

- Allison, R. (2008). Russia resurgent? Moscow's campaign to 'coerce Georgia to peace.' The Royal Institute of International Affairs. Retrieved October 12, 2022, from https://commonweb.unifr.ch/artsdean/pub/gestens/f/as/files/4760/39349_201918.pdf
- Bastos, M., & Mercea, D. (2018). The public accountability of social platforms: lessons from a study on bots and trolls in the Brexit campaign. The Royal Society Publishing. Retrieved October 12, 2022, from <https://doi.org/10.1098/rsta.2018.0003>
- Brodersen, K. H., Gallusser, F., Koehler, J., Remy, N., & Scott, S. L. (2015, January 1). Inferring causal impact using Bayesian structural time-series models. Google Research. Retrieved October 12, 2022, from <https://research.google/pubs/pub41854/>
- Campello, R.J.G.B., Moulavi, D., Sander, J. (2013). Density-Based Clustering Based on Hierarchical Density Estimates. Advances in Knowledge Discovery and Data Mining. PAKDD 2013. Lecture Notes in Computer Science, vol 7819. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-642-37456-2_14
- Cunningham, C. (2020, November 18). A russian federation information warfare primer. The Henry M. Jackson School of International Studies. Retrieved October 12, 2022, from <https://jsis.washington.edu/news/a-russian-federation-information-warfare-primer/>
- Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019, May 24). Bert: Pre-training of deep bidirectional Transformers for language understanding. arXiv.org. Retrieved October 12, 2022, from <https://arxiv.org/abs/1810.04805>
- Funaiole, M., Bermudez J., Conley H., & Newlin C. (2021, April 22). Unpacking the Russian troop buildup along Ukraine's border. Center for Strategic and International Studies. Retrieve October 12,2022, from <https://www.csis.org/analysis/unpacking-russian-troop-buildup-along-ukraines-border>
- Grootendorst, M. (2022, March 11). Bertopic: Neural topic modeling with a class-based TF-IDF procedure. arXiv.org. Retrieved October 12, 2022, from <https://arxiv.org/abs/2203.05794>
- Hiatt, F. (1995, January 29). Truth of chechen war derails Russian Propaganda Machine. The Washington Post. Retrieved October 12, 2022, from <https://www.washingtonpost.com/archive/politics/1995/01/29/truth-of-chechen-war-derails-russian-propaganda-machine/62c813f3-9f8d-47e1-bed1-254443ab3d69/>
- King, B., & Jordan, D. (2022, June 27). Russia in debt default as payment deadline passes. BBC News. Retrieved October 12, 2022, from <https://www.bbc.com/news/business-61929926>
- National Intelligence Council (2021). Foreign Threats to the 2020 US Federal Election. Retrieved October 12, 2022, from <https://www.dni.gov/files/ODNI/documents/assessments/ICA-declass-16MAR21.pdf>
- Schafer, B. (2019, September 3). Hamilton 2.0 methodology & faqs. Alliance For Securing Democracy. Retrieved October 12, 2022, from <https://securingdemocracy.gmfus.org/hamilton-2-0-methodology-faqs/>
- Siebold, S., & Blenkinsop, P. (2022, February 27). EU tightens Russian sanctions and buys weapons for Ukraine. Reuters. Retrieved October 12, 2022, from <https://www.reuters.com/world/europe/eu-close-air-space-russia-curb-media-target-belarus-2022-02-27/>
- Wintour, P. (2022, January 20). Blinken to reassert Ukraine sovereignty in Berlin speech. The Guardian. Retrieved October 12, 2022, from <https://www.theguardian.com/us-news/2022/jan/20/antony-blinken-us-ukraine-russia-berlin-speech>
- Zhang, Z. (2022, March 30). Study Confirms Influence of Russian Internet "Trolls" on 2016 Election. Columbia SIPA. Retrieved October 12, 2022, from <https://www.sipa.columbia.edu/news/study-confirms-influence-russian-internet-trolls-2016-election>