Intersect, Vol 14, No 3 (2021)

Behind the COVID Curtain: Analyzing Russia's COVID-19 Response on Twitter Using Natural Language Processing and Deep Learning

Tal Feldman
Wake Forest University

This paper analyzes the Twitter activities of five Russian political institutions, in their source languages, to assess the country's communication response to the beginning of the COVID-19 pandemic. This study employs several Natural Language Processing techniques in English and Russian, including a deep learning model to determine tweet sentiment and a lexicon-based method to classify tweets as health related. I argue that there was a coordinated response across the different institutions. I find that most accounts in this study tweeted less during the pandemic. Further, I argue that the proportion of health-related tweets over time was guided more by political motives than health concerns. Finally, I observe a significant difference in sentiment across English and Russian tweets, with the English being more positive and the Russian being more neutral.

Introduction

Much can be said about a country's response to a crisis by analyzing official state messages. For countries like the Russian Federation, there is often little other information. In recent years, many political institutions have turned to social media, such as Twitter, to communicate with citizens and the world—official accounts of Presidents, Parliaments, and Prime Ministers from a multitude of countries can be found simply by going to Twitter's website. During the COVID-19 pandemic, colloquially known as the "Coronavirus," online communication became even more important. This study is one of the first to examine the Russian Federation's official response to COVID-19 by utilizing social media data, specifically Twitter.

Twitter posts, called "tweets," in English have been widely studied in the social science literature, especially political science (Curini & Franzese, 2020, p. 1068). However, much less has been written about Twitter in Russian. Although VKontakte (VK) is the most popular social media site in Russia, Twitter is still widely used, and is ranked in the top ten most popular internet resources in the country. While VK is used mostly by Russian-speakers, Twitter is used across the world (Smetanin, 2020). As a result, several political entities in this study, such as the President of Russia, have Twitter accounts both in English and Russian. Thus, Twitter was chosen for this study to investigate Russia's communications with the world, rather than just Russians.

The official Twitter accounts of five Russian political institutions—not the people heading these institutions—were selected for this study: the President of Russia, the Government of Russia, the State Duma, the Federation Council, and United Russia. These five institutions are some of the most important in the Russian Federation. At the center of Russia's political system is the President of Russia, currently Vladimir Putin. The Government of Russia (hereafter, "Government") exercises executive power and is composed of the Prime Minister, Deputy Prime Ministers, and Federal Ministers, such as the Minister of Foreign Affairs. The Government is appointed by the President. The State Duma is the lower house of Russia's parliament, composed of 450 elected members, and can be dissolved by the President (The Constitution of the Russian Federation, 1993). The Federation Council is the upper house of Russia's Parliament, comprising two appointed members from each of Russia's territorial subjects, totaling 170. Because Federation Council members are appointed by regional leaders, the chamber is largely viewed as dependent on the President (Remington, 2003, pp. 670–673). The last account selected for this study is not of a constitutionally created entity but the ruling political party, United Russia. United Russia, of which President Putin is a member, dominates Russia's politics—at the time of writing, United Russia members held 75 percent of seats in the Federation Council, 75 percent of the State Duma, 78 percent of seats in regional parliaments, and 88 percent of governorships (United Russia, 2020).

Tweets before and during COVID-19 were extracted from these accounts and examined using several Natural Language Processing (NLP) methods. NLP is a branch of computer science that aims at transforming raw text into data that can be analyzed by a computer. For example, identifying keywords and then searching text for those keywords is a simple NLP task (Kaur, 2014, pp. 168–169). I employ a similar method to determine whether tweets were related to health or not. Another task in NLP is sentiment analysis, in which the computer assigns text to sentiment classes, such as "positive" or "negative." To perform sentiment analysis, I use a more complex NLP tool, a deep learning model called BERT.

Deep learning is a subfield of artificial intelligence (AI) that utilizes artificial neural networks.[1] Deep learning models have proved to be extremely effective in complex NLP tasks, especially sentiment analysis

[1] Neural networks are AI models which are loosely based on biological neural systems. A neural network contains many interconnected processors, called neurons, which have weights that predict an outcome based on inputs. To train these weights, the neural network is given a training dataset from which it learns to map inputs to an output. A deep neural network is one in which there are many layers of neurons. In a deep neural network, or deep learning model, the input data is processed by neurons in the first layer, which then pass their weighted outputs to the neurons in the second layer, and so on. These weights are analogous to coefficients in a multivariate statistical model. However, unlike traditional statistical models, a deep neural network can have millions of weights.

(Habimana et al., 2019, p. 6). Sentiment analysis has been widely applied for English but remains underutilized for the Russian language (Smetanin, 2020, p. 110693). This study presents one of the first scholarly applications of the cutting-edge BERT deep learning model for sentiment analysis in Russian politics.

Having established this background, it is important to review the political climate during the time of this study. Before the brunt of the COVID-19 pandemic, the Russian political sphere was changing quickly. In January 2020, President Vladimir Putin announced major constitutional changes were to come. Subsequently, Prime Minister Dmitry Medvedev's Government resigned, and Putin appointed Mikhail Mishustin as the new Prime Minister. Getting close to his constitutionally mandated term limit in 2024, Putin began working on constitutional changes that would allow him to remain President until at least 2036 (Standish & Mackinnon, 2020).

As the virus began to spread, early official statistics of COVID-19 cases showed the virus was almost a nonissue in Russia. However, as the virus got worse, instead of making unpopular decisions, such as lockdowns, Putin pushed responsibility to other leaders, including regional ones ("As COVID-19 Spreads, Vladimir Putin Is Invisible," 2020). This was unexpected, because Putin has spent his time in power dismantling the asymmetric federalist system in favor of a more centralized one, taking steps to have tighter control over the regional governorships (Busygina et al., 2018, pp. 66–67). Although official Russian statistics maintained that virus cases and mortality were low, Putin's approval ratings dropped to historic lows ("Russia's Covid-19 Outbreak Is Far Worse than the Kremlin Admits," 2020). Still, Putin's constitutional amendments passed the referendum in July with a comfortable majority.

This study has two main contributions to the field: the first is the application of an assortment of cutting-edge computational techniques that will hopefully provide a framework for further research; the second is several new insights about Russia's COVID-19 communications response that came from utilizing these techniques. These new insights from the data support three main arguments. The first is that these Twitter accounts acted in a coordinated manner, perhaps at the instruction of the President. The second is that many of the phenomena observed in the data can be attributed to political event and motives, rather than the state of the health crisis. The third argument is that the English and Russian Twitter accounts of these institutions are fundamentally different from one another in the way they present current political situations.

Materials and methods

The data in this study was extracted directly from Twitter using the Twitter Application Programming Interface (API). Eight Twitter accounts, representing five Russian political entities, were chosen. Each of the five entities had an account in Russian and three had a second account in English, supposedly as a translation to the English-speaking world. I extracted all tweets from each of these accounts between September 7, 2019 and June 26, 2020.

These dates comprise a forty-two-week date range, centered on February 1, 2020, the date of the first COVID-19 case in the Russian Federation, according to the European Center for Disease Prevention and Control (2020). The latter boundary of the date range, twenty-one weeks from February 1 (hereafter, "during COVID-19"), was chosen because it marked the last week before the constitutional referendum, by which time much of Russia's political establishment had stopped focusing on the pandemic. To get a full sense of the influence of these early days of COVID-19, the former boundary of the date range was chosen to be twenty-one weeks in the opposite direction (hereafter, "before COVID-19"), well before it was a global pandemic.[2]

All tweets from each account over this date range were extracted, comprising a total of 8,602 tweets. This data included the text of each tweet, the date and time of its posting (converted from UTC to Moscow Time), and any hashtags used.[3] The name of each account, its primary language, which entity it represents, the number of followers it has, and the number of tweets that were extracted are provided in Table 1 for reference.

Table 1: The Twitter Accounts Analyzed in this Study

Account Name	Language	Institution	Followers Count (in thousands)[4]	Tweets Extracted
KremlinRussia	Russian	President of Russia	3,400	922
KremlinRussia_E	English	President of Russia	826	736
Pravitelstvo_RF	Russian	Government of Russia	617	1,835
GovernmentRF	English	Government of Russia	223	321
dumagovru	Russian	State Duma	922	1,473
state_duma	English	State Duma	2	903
SovFedInfo	Russian	Federation Council	149	801
er_novosti	Russian	United Russia	156	1,611

As evidenced by Table 1, the President of Russia, the Government, and the State Duma had second Twitter accounts in English, while the Federation Council and United Russia did not. The number of followers for each account varied widely, with KremlinRussia amassing over three

[2] Instead of going back twenty-one weeks from February 1, I considered using the same dates but from 2019 (February 1, 2019 to June 26, 2019). This might help minimize the disparities among tweeting frequency associated with annual events, such as holidays. However, I decided against this because this date range would lose the perspective of political issues that started before the pandemic and continued during it.

[3] The number of likes, comments, and retweets of each post were not used in this study due to the wide date range.

[4] The number of followers for each account was extracted on July 6, 2020. They are estimated to the nearest thousand.

million followers and state_duma attracting around two thousand. The number of tweets also varied between different accounts with the most tweets coming from Pravitelstvo_RF and the least from its English counterpart, GovernmentRF.

All subsequent text analysis was conducted in the original language of the data. While it would have been possible to translate all of the tweets into one language, this method would have minimized the credibility of this research. Although translating tools such as Google Translate have made impressive strides in the last several years, using source language models for sentiment analysis is likely to yield more accurate results (Balahur & Turchi, 2014, p. 69).

The most basic examination performed in this paper is that of tweet frequency across different accounts. Beyond this, I implemented two methods to analyze the text of the tweets. The first classifies tweets by whether they are health related or not. The second evaluates the sentiment of each tweet using a deep learning model.

Classifying tweets as health-related: A lexicon-based approach

It was imperative to examine how much the Twitter accounts were tweeting about health. To do this, I created numerous Regular Expression (regex) search patterns, both in English and Russian, that were related to health and the pandemic. If a tweet contained any of these search patterns, ignoring case, it was tagged as related to health.

Although topic modeling is a rich area of NLP, simple regex methods have outperformed "carefully designed and incrementally improved DNN-based supervised-learning classifiers" for COVID-19 classification (Markov et al., 2021). Indeed, regex approaches are particularly well suited for tasks with small sets of keywords and phrases, such as for COVID-19 (Markov et al., 2021). This kind of regex keyword model to extract tweets that are related to COVID-19 has already been implemented, using keywords such as "coronavirus" and "covid" (Müller et al., in press, p. 2). This paper broadens this categorization of tweets to encompass anything health related, while focusing on words related to the pandemic. For instance, a tweet mentioning doctors, masks, or quarantine during the time period in question might not explicitly contain the word "Coronavirus," but it is still useful for this study. Hence, the broader classification system used in this study cannot precisely say if a tweet was related to COVID-19, but most health-related tweets likely were, implicitly or explicitly.

I chose thirty-three regex search patterns in English and thirty-eight in Russian.[5] Much consideration was put into picking search patterns that would capture different forms of the same words. For example, in English, the root "epidem" was selected, because it would capture "epidemic," "epidemiologist," "epidemiologists," and more. For words that could be

[5] The regex search patterns were compiled by the author, who is fluent in Russian, and reviewed by several other native Russian speakers.

contained within other words, I specified that no characters could precede them, follow them, or both.

A limitation of this method is that it is not exact. Despite careful selection, it is possible that some tweets were erroneously labeled as health-related and some that should have been labeled were not. However, topic classification is an evolving field with no perfect solutions. Still, other studies have used lexicon-based approaches like this one to identify health-related texts in social media data (Park & Conway, 2018).

Sentiment analysis: Introducing BERT

Building on classifying tweets as health-related, I analyzed the sentiment of the Twitter data before and during COVID-19. I employed two models based on the cutting-edge bidirectional unsupervised deep learning model, Bidirectional Encoder Representations from Transformers, or BERT.

Sentiment analysis has long been an important task in NLP and has seen many possible solutions over the years. The lexicon-based models use lexicons with predefined sentiments for certain words to classify text into sentiment classes—such as "positive," "neutral," and "negative." SentiStrength, a lexicon-based sentiment algorithm available in Russian (Thelwall et al., 2010, p. 2549), has been used to measure sentiment in Russian political discourse on Twitter (Spaiser et al., 2017, p. 137). However, even better results have been achieved with deep learning models (Habimana et al., 2019, pp. 3–4).

The last decade saw major advances in applying deep learning models to NLP.[6] However, perhaps the biggest recent advancement is Google's BERT, which achieved state-of-the-art results on a wide variety of NLP tasks. BERT models are first pre-trained on a large unlabeled corpus and then fine-tuned on labeled data for more specific tasks (Devlin et al., 2019, pp. 4172–4175).

There are relatively few accurate Russian sentiment analysis models, and large publicly available training datasets are scarce (Smetanin, 2020, p. 110709). This is why the open-sourcing of the BERT multilingual model, which has been pre-trained for over one hundred languages on massive corpuses, was an exciting advancement for NLP in many languages, including Russian.[7] Because sentiment analysis is best performed in the original language, this study uses two different BERT-based models for sentiment analysis, one for tweets in Russian and the other for tweets in English. Both models were fine-tuned at the Moscow Institute of Physics

[6] The skip-gram model was used to create word2vec. Combining global factorization and local context window methods, such as in the skip-gram model, led to the popular GloVe model (Habimana et al., 2019, p. 6). Another breakthrough came in the form of ELMo, a model that introduced deep contextualized representations for existing neural architectures (Peters et al., 2018, p. 2232).

[7] These open-source models can be found at https://github.com/google-research/bert. For the multilingual models, details of the currently-available models can be found at https://github.com/google-research/bert/blob/master/multilingual.md.

and Technology and released on DeepPavlov, a popular open-source NLP library (Burtsev et al., 2018, p. 122).

Both models have twelve layers and over one hundred million parameters. The Russian BERT model was initialized on the multilingual version of BERT and further trained on Russian Wikipedia, news data, and social media posts. The model achieved an F-1 score of 77.42 after fine-tuning on the 3-class RuSentiment[8] (DeepPavlov, 2018), an increase over previous neural network models (Kuratov & Arkhipov, 2019, pp. 334–336). The English BERT model was pre-trained on the BooksCorpus and English Wikipedia, which sum up to over three billion words (Devlin et al., 2019, pp. 4172–4175). It was further trained on data from Twitter, Reddit, and Facebook News comments. The model was fine-tuned on the fine-grained Stanford Sentiment Treebank, SST-5 (DeepPavlov, 2019). SST-5 is composed of movie reviews categorized into five classes—"very negative," "negative," "neutral," "positive," and "very positive." Naturally, models perform much better on classifying the binarized corpus[9] than SST-5 (Socher et al., 2013, pp. 1631–1633). However, this study uses a model with the more robust, 5-class approach whose test accuracy was 67.15 percent on SST-5 (DeepPavlov, 2018).

As neural models require little pre-processing of text data, I excluded URLs but kept hashtags, following procedure in other sentiment analysis studies using neural networks (Rogers et al., 2018, p. 756). I applied the English and Russian BERT models on the English and Russian tweets, respectively. A limitation of using two models is that they are trained on different data in different languages. Although both are pre-trained and fine-tuned on similar datasets of social media and news posts, these models are inherently different. This is advantageous because of the vast differences between English and Russian, including the way in which sentiment is expressed in these two languages, but this is important to note when comparing the results of the two models. An alternative would have been to translate all tweets into one language, but text can lose many important characteristics in translation that hinder effective sentiment analysis. Despite the improvements in statistical machine translation, there is still a gap in sentiment classification performance in favor of models trained on source language data (Balahur & Turchi, 2014, p. 69).

Results and analysis

To contextualize the severity of the pandemic during the time studied here, I created the graph in Figure 1 that depicts the number of reported daily cases in the Russian Federation from December 31, 2019—the date of the

[8] RuSentiment is one of the most popular Russian sentiment datasets. It is composed of VK posts in Russian, labeled on a three-class scale—"negative," "neutral," and "positive." There is also a "speech acts" class, constituting expressions of gratitude and congratulations, that is a subclass of the "positive" class.

[9] SST-5 can be binarized into simply positive and negative by removing the "neutral" class and merging "very negative" with "negative" and "very positive" with "positive."

first reported case of "viral pneumonia" in Wuhan City (World Health Organization, 2020)—to June 26, 2020. The solid black line represents the number of daily cases and the dashed red line marks the first reported COVID-19 case in Russia. Figure 1 is based on the data reported by the European CDC (2020) and helps to inform the official messages of political institutions about the severity of the pandemic.

Figure 1

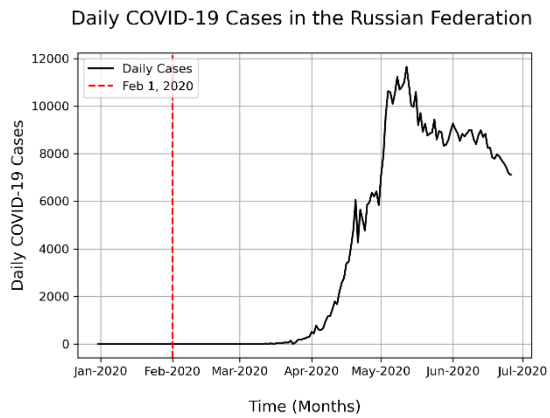

The results are split into three sections. The first section presents the tweeting frequencies among the different accounts, in which I argue that the discrepancies in frequencies are due to the heightened responsibility of certain institutions during COVID-19. The second section examines the proportions of tweets related to health, in which I argue that the changes in health-related tweets across most accounts can perhaps be attributed more to political motives and events than the severity of the health crisis. The third section provides an analysis of sentiment, in which I build on the previous sections to argue that the English and Russian accounts have fundamentally different objectives and behaviors.

Tweeting frequency: Before and during COVID-19

First, I aggregated each account's tweets into weekly frequencies (defined as seven-day periods). This aggregation avoids measuring discrepancies among tweeting habits during different days of the week, such as weekdays and the weekends. Subsequently, I analyzed the difference in number of weekly tweets before and after the first reported Russian case of COVID-19. Splitting the dataset into pre and during COVID-19 produced two subsets that each contained data over twenty-one weeks, with the former subset serving as a baseline to analyze alterations in tweeting habits of these accounts before and during COVID-19. The difference in means for each Twitter account was calculated using a one-tailed t-test. These results are presented in Table 2.

Table 2: T-test Results—Weekly Tweeting Frequencies Before and During COVID-19

Account	Mean of Weekly Tweet Frequency pre-COVID-19 ($\bar{x}_1$)	Mean of Weekly Tweet Frequency during COVID-19 ($\bar{x}_2$)	Difference in Means ($\bar{x}_1 - \bar{x}_2$)
KremlinRussia	26.6 (3.0)	17.3 (0.4)	9.3** (0.6)
KremlinRussia_E	21.0 (2.4)	14.0 (0.5)	7.0** (0.5)
Pravitelstvo_RF	41.5 (2.9)	45.9 (5.4)	−4.3 (0.9)
GovernmentRF	10.5 (2.2)	4.8 (1.2)	5.7** (0.9)
dumagovru	42.0 (4.1)	28.1 (2.6)	14.0** (0.8)
state_duma	24.8 (2.0)	18.2 (0.9)	6.5** (0.4)
SovFedInfo	20.0 (2.2)	18.2 (1.5)	1.8 (0.6)
er_novosti	29.2 (2.5)	47.5 (4.1)	−18.2*** (0.7)
Notes: Standard errors in parentheses. ***$p < 0.001$, ** $p < 0.01$, * $p < 0.05$			

Table 2 shows that most of the accounts tweeted less per week during COVID-19, with five of them tweeting significantly less ($\alpha = 0.05$). These were both accounts of the State Duma, both of the President's accounts, and the English account of the Government. Amidst the international crisis, some of these accounts changed their tweeting patterns drastically. KremlinRussia and dumagovru tweeted over thirty percent less, while GovernmentRF tweeted more than fifty percent less. SovFedInfo, the Federation Council's account, tweeted less, but not significantly so. Perhaps there were fewer special events to tweet about—such as state visits at the Kremlin or bills under consideration by the Duma—but with so many challenges presented by the pandemic, it is difficult to attribute this suddenly more silent approach to a lack of topics to tweet about.

Still, there were two accounts which tweeted more. Pravitelstvo_RF, the Government's Russian account, tweeted more, although not significantly. United Russia's account, er_novosti, on the other hand, tweeted significantly more—over sixty percent more. While the rest of the political establishment decreased tweeting frequency, these accounts went in the opposite direction. I argue that this split in account activity, with some tweeting significantly more and others tweeting significantly less, is the result of President Putin devolving some responsibility from his office during the pandemic to institutions such as the Government and local leaders.

During the pandemic, perhaps to avoid responsibility for decisions with negative consequences, President Putin left decisions on how to deal with the crisis up to local leaders, such as governors and mayors, as mentioned above. Mayor Sobyanin of Moscow, for example, earned notoriety for his strong response to the pandemic, even telling President Putin in a televised conference that official figures of the pandemic were not accurate due to the lack of testing. It is likely, however, that this display from the Mayor, a member of Putin's United Russia, was a coordinated effort (Malpas & AFP, 2020).

Thus, while almost all accounts tweeted less during the pandemic, the two that did not were also more important in Russia's pandemic response. The Government took place as a more central figure in the pandemic response, which is reflected in their increase in tweets. Furthermore, during this time, the Government appeared to focus more on its domestic audience than international one: it tweeted far more in Russian. United Russia, the party of most governors and regional parliaments, finding its regional leaders center stage in the fight against COVID-19, tweeted significantly more during the pandemic. Meanwhile, some central institutions, such as the President and parliament, tweeted less, perhaps backing away from the spotlight during this difficult time.

Looking at the trends more broadly, it is evident that all the English accounts tweeted dramatically less than their Russian counterparts. This suggests different messages for the Russian and English audiences. This will be expounded on in the next two sections.

Tweeting about health

This section presents results from the analysis of health-related tweets. Figure 2 presents the results of the lexicon-based approach to tagging tweets as health-related. The data was aggregated by week, like before, and then the number of health-related tweets for each week was summed and divided by the overall number of tweets that week to calculate the proportion of health-related tweets. The proportions of health-related tweets of Russian accounts are plotted in blue. The proportions for the three English accounts are plotted in green with their Russian counterparts for easier examination. The date range for Figure 2 is the same as in Figure 1, and Figure 2 also includes the dashed red line to denote the first reported case in Russia.

Figure 2: Weekly Proportions of Health-Related Tweets Across Twitter Accounts

Graph 1
Weekly Proportion of Health-Related Tweets from the President of Russia

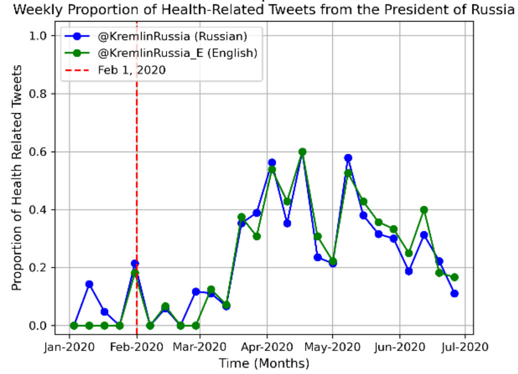

Graph 2
Weekly Proportion of Health-Related Tweets from the Government of Russia

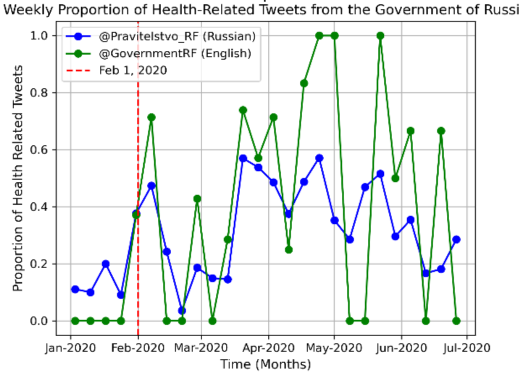

Graph 3
Weekly Proportion of Health-Related Tweets from the State Duma

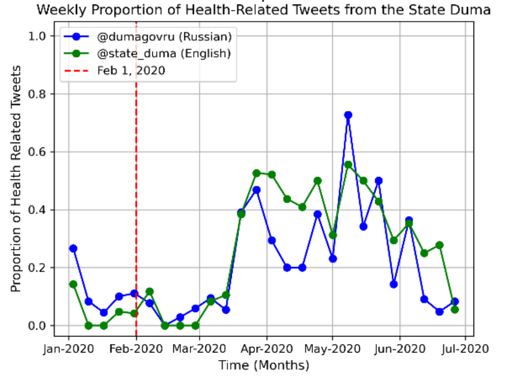

Graph 4
Weekly Proportion of Health-Related Tweets from the Federation Council

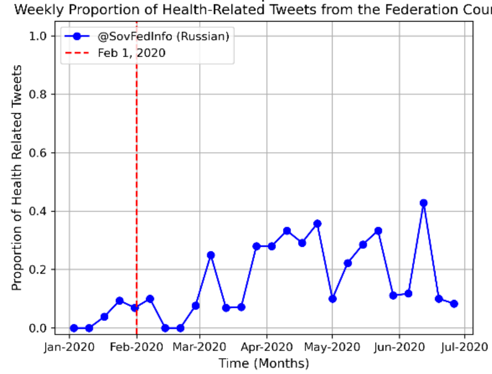

Graph 5
Weekly Proportion of Health-Related Tweets from United Russia

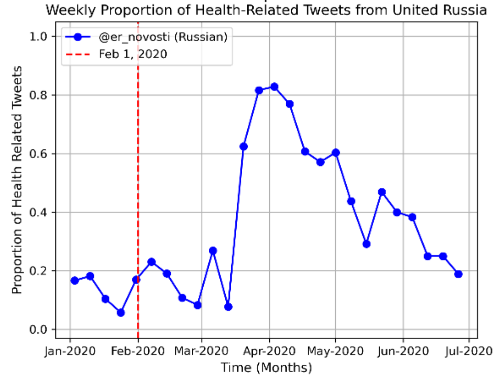

Figure 2 strengthens this paper's argument that these proportions were guided more by political motives than the health crisis itself. When the first case was reported in February, the state's message was that the virus was under control. All of the Twitter accounts had a spike in health-related tweets at this time—perhaps because Russia closed parts of its border with China (Charles Maynes, 2020). After this initial jump, all of the proportions decreased, suggesting that health was not at the forefront of conversation, even though the pandemic was growing in strength across the world. At the end of February, Putin announced the date for the constitutional referendum: April 22 (Jack Arnholz, 2020). With national attention on the upcoming referendum, proportions of health tweets remained very low, as did the number of reported cases. In the first weeks of March, both the State Duma and Federation Council approved the constitutional amendments (Andrew Osborn & Polina Ivanova, 2020), followed by the Constitutional Court (Higgins, 2020a).

Slowly, each account increased their health-related tweeting proportions. This rise is noticeable in most accounts between the second and third week of March. At this time, Russia took strong countermeasures against COVID-19: creating the Coordination Council to control the spread of the virus, shutting down its borders, allocating hundreds of billions of rubles for economic stimulus, and more (Mikhail Mishustin, 2020). Moscow's Mayor Sobyanin also asked Muscovites to work from home. These sweeping changes came suddenly—in a matter of several days—and are clearly reflected in the data (Standish, 2020). By the end of March, Putin postponed the referendum (Higgins, 2020b). From there, the proportions climbed.

By the middle of April, state media and powerful political leaders, including the President, acknowledged the severity of the virus (Troianovski, 2020). Most of the accounts reached a peak of health-related tweets in April and early May (SovFedInfo being the exception, although its proportions remained relatively low throughout). This increase in proportions corresponds to the rising number of daily cases over the same period in Figure 1. However, the peak also corresponds with Russians' growing agitation.

Independent polling by the Levada Center from the end of April showed that more people disapproved of Putin and the Government's response to the pandemic than that of mayors and governors ("Putin's Virus Response Earns Lower Marks Than Local Leaders'," 2020). In early May, Putin's approval rating hit a historic low, dipping below 60 percent ("Putin's Approval Rating Drops to Historic Low," 2020). Subsequently, on May 11, Putin declared the end of non-working days. Perhaps not coincidentally, May 11 happened to mark the peak of daily reported cases (TASS, 2020). After this date, daily cases started decreasing, and the accounts were already in full retreat from tweeting about health.

By late May, Putin turned attention to the belated Victory Day Parade (Al Jazeera, 2020). By the middle of June, the Kremlin maintained that the virus was under control, and Moscow lifted its remaining restrictions, just in time for the Parade and referendum (Ellyatt, 2020). Although the number of daily cases was still high towards the end of the time period in this study, most of the accounts tweeted about health less than twenty percent of the time. Indeed, even though there were reportedly over 7,000 daily cases at the time, the proportions of health-related tweets returned to roughly the same numbers as on February 1, 2020. With the official messaging away from the health crisis, the constitutional referendum, an important exercise for ascertaining Putin's legitimacy, passed comfortably, reportedly with seventy-eight percent approving the measures (Higgins, 2020c). Even if the numbers in Figure 1 are mostly accurate, the behavior of these accounts still suggests that political motives were the driving factor, not the health crisis itself. The similarity in behavior across these different political institutions suggests a unified public relations effort, albeit with minor variations.

Looking just at the proportions for Russian accounts, it is evident that United Russia and the Government of Russia tweeted more about health than the rest. The Government had eight weeks during which health was the topic of more than forty percent of tweets, while United Russia had ten such weeks, adding to the argument that these institutions had heightened importance during COVID-19. The President of Russia, on the other hand, had only three such weeks, while the Federation Council had only one. Nevertheless, most accounts' proportions were low overall, considering the severity of the health crisis, suggesting the official message remained one of control of the virus.[10]

Looking at Graphs 1-3, some more insights about the differences across the accounts in English and Russian emerge. All three English graphs roughly follow the trend outlined above. (GovernmentRF looks more sporadic, but this can be attributed to its low number of weekly tweets during this period, as evidenced in Table 2). I calculated the differences and found that the overall proportions of health-related tweets from all of the English accounts were higher than their Russian counterparts during COVID-19.

Performing a chi-square test of independence, in which the null hypothesis is that there is no relationship between the variables,[11] I found a significant relationship between the language of the account and proportion of health-related tweets for the Government and State Duma accounts during COVID-19 ($\alpha = 0.01$). The following are the results of the chi-square test for the Government and State Duma accounts during COVID-19, respectively: $\chi^2(1, N = 1064) = 14.7,\ p < 0.01$ and $\chi^2(1, N = 973) = 10.1,\ p < 0.01$. Therefore, the proportions of health-related tweets from these two English accounts were significantly different from their Russian counterparts. This builds on the differences presented in Table 2—the accounts geared at the English-speaking world tweeted far less during COVID-19, but a higher percentage of their tweets were related to health. This suggests that to the outside world, the accounts were projecting an image of health being a high priority, while frequently tweeting about other topics to the Russian people.

Sentiment analysis

This section presents the results of the sentiment analysis by the two BERT models. A very clear tendency emerged: according to the BERT models, the Russian tweets were mostly neutral, while the English tweets were mostly very positive.[12] The differences were staggering: over 98 percent of tweets from Russian accounts were classified as "neutral" before and during COVID-19. Accordingly, roughly 1 percent of Russian tweets were classified as "positive." On the other hand, around 66 percent of English tweets were "very positive" before COVID-19 and 58 percent were "very positive" during COVID-19. Moreover, the percentages of "negative" and "very negative" tweets in both languages were extremely low. As

[10] The proportion of United Russia's tweets reached above 0.8, but quickly retreated. The rest of the accounts' proportions remained much lower.
[11] The alternative hypothesis is that a relationship between the variables exists. I assumed independence of samples for this test due to the vast difference in weekly tweets presented in Table 2. (A z-test for two proportions produces the same p-value as a chi-square test).
[12] Recall that the English model categorized data into five classes, while the Russian categorized into three.

mentioned earlier, there are inherent problems with using two different models, but these differences are overwhelming.

Looking closer at the data, the sentiments of the Russian accounts did not change much before and during COVID-19, with all accounts remaining over 96 percent "neutral" throughout both time periods. However, the proportions of "very positive" tweets in English did change across the three accounts. The Government's English tweets went from 84 percent "very positive" before COVID-19 to 62 percent during COVID-19. The Duma's retreated from 64 percent "very positive" to 58 percent. The President of Russia's English tweets took a smaller dip, going from 59 percent "very positive" to 57 percent. These differences are not significant when taking the possible error of the model into account, but they are still important to note.

To assess the overall differences between Russian and English accounts, I performed a test for the difference in proportions over the entire time range. I assumed the errors of the model were evenly distributed among the sentiment classes. This assumption most likely overestimates error, making the p-value higher, because errors in classification "cancel each other out." For instance, if the model incorrectly classified a negative text as neutral and a neutral text as negative, the number of neutral and negative texts would still be correct. However, due to the magnitude of the difference, even with this assumption, the results of the test are significant. Furthermore, it was necessary to merge sentiment classes to compare the 5-class English model to the 3-class Russian model.[13] The results of the test are significant ($\alpha = 0.001$), as seen in Table 3.

Table 3: Difference in Proportions of Sentiment Classes Across English and Russian Accounts

[13] When binarizing the English model, as noted earlier, the "neutral" class is removed and the two categories on either side are merged (Socher et al., 2013, pp. 1631–1633). This method combines "very positive" with "positive" to yield an aggregate positive category and combines "very negative" with "negative" to yield an aggregate negative category. I followed this same procedure to transform the 5-class English results into 3 classes, preserving the original "neutral" category and combining the other categories, as in the binarization method.

Sentiment Class	Russian Proportion ($\hat{p}_1$)	English Proportion ($\hat{p}_2$)	Difference in Proportions ($\hat{p}_1 - \hat{p}_2$)
"Neutral"	0.988 (0.075)	0.300 (0.066)	0.630*** (0.100)
"Positive"	0.011 (0.075)	0.628 (0.131)	−0.617*** (0.151)
Notes: Errors in parentheses. There was no significant difference among "negative" proportions. ***p < 0.001, ** p < 0.01, * p < 0.05			

Table 3 sheds light on just how different the English and Russian accounts are. Perhaps the Russian accounts address similar topics as the English ones, but the way in which topics are addressed is drastically different across the two languages. The sentiment in messages for English-speaking audiences is significantly more positive and less neutral than for Russian-speaking ones.[14]

This presents new findings about the differences among these accounts. The previous sections highlighted that the number of tweets of the English accounts were far lower than that of their counterparts, suggesting that they were more selective in choosing English tweet topics. It was also proven that most of these accounts tweeted proportionally more about health. Table 3 furthers these results by showing the stark differences in sentiment classification across the two languages. These results solidify the impression that the English accounts are, in fact, not mere translations of the Russian ones. The accounts behave very differently and likely have different purposes. The English and Russian accounts tweet at different frequencies, about different topics, and with starkly different sentiments. These differences beckon the following conclusion: the English accounts are a postured version of what is conveyed to the Russian people, aimed to impress the English-speaking world.

Conclusion

This article presented one of the first analyses of Russia's communications response to the COVID-19 health crisis using social media data. Moreover,

[14] For example, in discussing COVID-19, the Government's English account kept a "very positive" attitude according to BERT, tweeting "On Russia Day, the Russian tricolour was projected onto the Kremlin, Government House and Manezh, as well as videos showing today's national heroes: doctors and all those involved in the efforts to battle the COVID-19 pandemic…" Similarly, the state_duma tweeted "#Bill A new package of measures to support citizens and businesses during the #coronavirus was adopted by the State Duma. The measures are aimed at supporting families with children, self-employed, small and medium-sized businesses, tourists, and many others." However, the accounts did not tweet only positively. The following tweet from state_duma was classified as "negative": "State Duma Commission will review fake news about allegedly inconsistent data on coronavirus mortality rate in #Russia published by @nytimes and @FT." Some of the President's English tweets were "negative," such as "The President announced the postponement of Victory Day celebrations due to the coronavirus epidemic."

this study is one of few to use deep learning for Russian political communication research. It found compelling evidence that much of the response was guided by political motives. Most accounts tweeted less, while some took center stage. For Russian audiences, the message was more neutral and proportionally less about health. For English-speaking audiences, there was a different narrative: a positive one in which Russia focused more on health issues. The markedly similar trends across Twitter accounts suggest a centralized response, perhaps guided in part or in whole by the President, modified for different audiences.

This study empirically corroborated some of what the news media had observed qualitatively about Russia's COVID-19 response. It used NLP and cutting-edge deep learning models to investigate unstructured text data. Further studies could use such advanced computational tools to analyze myriad Russian political communications; with the sinking popularity of central political leaders, research efforts might be well targeted at examining social media communications among Russian citizens about leaders' response to COVID-19 and their implications for the current regime.

The research techniques used here open possibilities for studies of communications in over one hundred languages. Accordingly, future studies could use these methods to investigate whether the trends found in this paper are unique to the Russian Federation during the pandemic or occurred elsewhere, too. With possible changes in the world order after COVID-19, this mode of research may prove effective in cross-national and multilingual comparisons and, perhaps, in predicting political trends across the world.

Acknowledgements

I thank the two anonymous reviewers for their insightful suggestions. I am also grateful to the attendees of the National Collegiate Research Conference at Harvard University and the London School of Economics Undergraduate Political Review Annual Conference for helpful comments. I am grateful to Faina Feldman for her invaluable support and Ashley Peake for her comments on the manuscript.

References

Al Jazeera. (2020, May 26). Coronavirus "past peak" in Russia, WWII parade in June: Putin. *Al Jazeera English*. https://www.aljazeera.com/news/2020/05/coronavirus-peak-russia-wwii-parade-june-putin-200526144211655.html

Andrew Osborn & Polina Ivanova. (2020, March 11). Sixteen more years? Russian parliament backs move to keep Putin in power. *Reuters*. https://www.reuters.com/article/us-russia-putin-idUSKBN20Y0VC

As COVID-19 Spreads, Vladimir Putin is Invisible. (2020, April 8). *The Economist*. http://www.economist.com/europe/2020/04/08/as-covid-19-spreads-vladimir-putin-is-invisible

Balahur, A., & Turchi, M. (2014). Comparative experiments using supervised learning and machine translation for multilingual sentiment analysis. *Computer Speech & Language*, *28*(1), 56–75. https://doi.org/10.1016/j.csl.2013.03.004

Burtsev, M., Seliverstov, A., Airapetyan, R., Arkhipov, M., Baymurzina, D., Bushkov, N., Gureenkova, O., Khakhulin, T., Kuratov, Y., Kuznetsov, D., Litinsky, A., Logacheva, V., Lymar, A., Malykh, V., Petrov, M., Polulyakh, V., Pugachev, L., Sorokin, A., Vikhreva, M., & Zaynutdinov, M. (2018). DeepPavlov: Open-Source Library for Dialogue Systems. *Proceedings of ACL 2018, System Demonstrations*, 122–127. https://doi.org/10.18653/v1/P18-4021

Busygina, I., Filippov, M., & Taukebaeva, E. (2018). To decentralize or to continue on the centralization track: The cases of authoritarian regimes in Russia and Kazakhstan. *Journal of Eurasian Studies*, *9*(1), 61–71. https://doi.org/10.1016/j.euras.2017.12.008

Charles Maynes. (2020, January 30). Russia Closes Parts of Border with China to Combat Spread of Virus. *Voice of America*. https://www.voanews.com/science-health/coronavirus-outbreak/russia-closes-parts-border-china-combat-spread-virus

Curini, L., & Franzese, R. (2020). Deep Learning for Political Science. In *The SAGE Handbook of Research Methods in Political Science and International Relations* (pp. 1053–1078). SAGE Publications. http://ebookcentral.proquest.com/lib/wfu/detail.action?docID=6178698

DeepPavlov. (2018). *Classification Models in DeepPavlov*. http://docs.deeppavlov.ai/en/master/features/models/classifiers.html#bert-models

DeepPavlov. (2019). *BERT in DeepPavlov*. http://docs.deeppavlov.ai/en/master/features/models/bert.html#id23

Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 4171–4186. https://doi.org/10.18653/v1/N19-1423

Ellyatt, H. (2020, June 11). Russia's coronavirus cases just passed 500,000—But it's easing lockdown fast. *CNBC*. https://www.cnbc.com/2020/06/11/russia-lifts-restrictions-ahead-of-victory-day-and-constitutional-vote.html

Habimana, O., Li, Y., Li, R., Gu, X., & Yu, G. (2019). Sentiment analysis using deep learning approaches: An overview. *SCIENCE CHINA Information Sciences*, *63*(1), 111102. https://doi.org/10.1007/S11432-018-9941-6

Higgins, A. (2020a, March 16). Russia's Highest Court Opens Way for Putin to Rule Until 2036. *The New York Times*. https://www.nytimes.com/2020/03/16/world/europe/russia-putin-president-for-life.html

Higgins, A. (2020b, March 25). Citing Virus, Putin Suspends Vote on Keeping Him in Power. *The New York Times*. https://www.nytimes.com/2020/03/25/world/europe/russia-putin-coronavirus.html

Higgins, A. (2020c, July 1). The Theatrical Method in Putin's Vote Madness. *The New York Times*. https://www.nytimes.com/2020/07/01/world/europe/putin-referendum-vote-russia.html

Jack Arnholz. (2020, February 27). Russia announces date for referendum that could extend Putin's rule. *ABC News*. https://abcnews.go.com/International/russia-announces-date-referendum-extend-putins-rule/story?id=69257041

Kaur, G. (2014). USAGE OF REGULAR EXPRESSIONS IN NLP. *IJRET: International Journal of Research in Engineering and Technology*, *3*(1), 168–174. https://doi.org/10.15623/ijret.2014.0301026

Kuratov, Y., & Arkhipov, M. (2019). Adaptation of Deep Bidirectional Multilingual Transformers for Russian Language. *Computational Linguistics and Intellectual Technologies*, *18*, 333–339. http://www.dialog-21.ru/media/4872/_-dialog2019scopus_rev2.pdf

Malpas, A. & AFP. (2020, March 28). Dour Moscow Mayor Comes to Fore as "PM for Coronavirus." *The Moscow Times*. https://www.themoscowtimes.com/2020/03/28/dour-moscow-mayor-comes-to-fore-as-pm-for-coronavirus-a69783

Markov, I. L., Liu, J., & Vagner, A. (2021). Regular Expressions for Fast-response COVID-19 Text Classification. *ArXiv:2102.09507 [Cs]*. http://arxiv.org/abs/2102.09507

Mikhail Mishustin. (2020, March 16). *Meeting with deputy prime ministers* [Press Release]. The Russian Government. http://government.ru/en/news/39161/#

Müller, M., Salathé, M., & Kummervold, P. E. (in press). *COVID-Twitter-BERT: A Natural Language Processing Model to Analyse COVID-19 Content on Twitter*. arXiv. http://arxiv.org/abs/2005.07503

Park, A., & Conway, M. (2018). Tracking Health Related Discussions on

Reddit for Public Health Applications. *AMIA Annual Symposium Proceedings*, *2017*, 1362–1371. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5977623/

Peters, M., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., & Zettlemoyer, L. (2018). Deep Contextualized Word Representations. *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, *1*, 2227–2237. https://doi.org/10.18653/v1/N18-1202

Putin's Approval Rating Drops to Historic Low: Poll. (2020, May 6). *The Moscow Times*. https://www.themoscowtimes.com/2020/05/06/putins-approval-rating-drops-to-historic-low-poll-a70199

Putin's Virus Response Earns Lower Marks Than Local Leaders. (2020, April 30). *The Moscow Times*. https://www.themoscowtimes.com/2020/04/30/putins-virus-response-rated-lower-than-local-leaders-poll-a70156

Remington, T. F. (2003). Majorities without Mandates: The Russian Federation Council since 2000. *Europe-Asia Studies*, *55*(5), 667–691. https://doi.org/10.1080/0966813032000086828

Rogers, A., Romanov, A., Rumshisky, A., Volkova, S., Gronas, M., & Gribov, A. (2018). RuSentiment: An Enriched Sentiment Analysis Dataset for Social Media in Russian. *Proceedings of the 27th International Conference on Computational Linguistics*, 755–763. https://www.aclweb.org/anthology/C18-1064

Russia's covid-19 outbreak is far worse than the Kremlin admits. (2020, May 21). *The Economist*. http://www.economist.com/europe/2020/05/21/russias-covid-19-outbreak-is-far-worse-than-the-kremlin-admits

Smetanin, S. (2020). The Applications of Sentiment Analysis for Russian Language Texts: Current Challenges and Future Perspectives. *IEEE Access*, *8*, 110693–110719. https://doi.org/10.1109/ACCESS.2020.3002215

Socher, R., Perelygin, A., Wu, J., Chuang, J., Manning, C. D., Ng, A., & Potts, C. (2013). Recursive Deep Models for Semantic Compositionality Over a Sentiment Treebank. *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, 1631–1642. https://www.aclweb.org/anthology/D13-1170

Spaiser, V., Chadefaux, T., Donnay, K., Russmann, F., & Helbing, D. (2017). Communication power struggles on social media: A case study of the 2011–12 Russian protests. *Journal of Information Technology & Politics*, *14*(2), 132–153. https://doi.org/10.1080/19331681.2017.1308288

Standish, R. (2020, March 18). Fortress Russia Begins to Show Cracks as Coronavirus Spreads. *Foreign Policy*. https://foreignpolicy.com/2020/03/18/fortress-russia-coronavirus-spread-covid-pandemic/

Standish, R., & Mackinnon, A. (2020, January 16). Who Is Russia's New

Prime Minister? *Foreign Policy*. https://foreignpolicy.com/2020/01/16/who-is-russias-new-prime-minister-mikhail-mishustin/

TASS. (2020, June 13). Coronavirus Epidemic in Russia Reached its Peak on May 11. *Russian News Agency TASS*. https://tass.com/russia/1167323

The Constitution of the Russian Federation, § Articles 95-111 (1993). http://www.constitution.ru/en/10003000-07.htm

Thelwall, M., Buckley, K., Paltoglou, G., Cai, D., & Kappas, A. (2010). Sentiment strength detection in short informal text. *Journal of the American Society for Information Science and Technology*, *61*(12), 2544–2558. https://doi.org/10.1002/asi.21416

Troianovski, A. (2020, April 13). Putin's Bleak Covid-19 Admission: 'We Don't Have Much to Brag About.' *The New York Times*. https://www.nytimes.com/2020/04/13/world/europe/coronavirus-russia-putin.html

United Russia. (2020). *Кто есть кто [Who is who]*. Единая Россия Официальный Сайт Партии [United Russia: Official Party Site]. https://er.ru/persons

World Health Organization. (2020, June 29). *Timeline of WHO's response to COVID-19* [Statement]. https://www.who.int/news-room/detail/29-06-2020-covidtimeline