



ROBOTROLLING

PREPARED AND PUBLISHED BY THE
**NATO STRATEGIC COMMUNICATIONS
CENTRE OF EXCELLENCE**



Executive Summary

In this edition of *Robotrolling* we track the most significant increase in inauthentic Russian-language social media activity we have observed since we began in 2017. While the level of bot activity remains much lower than four years ago, the uptick is concerning.

The increased activity coincided with the spring and summer military exercise season, and the period running up to the Russian Federation's Zapad exercises, scheduled for September 2021.

While fake activity increased in the Russian-language space, we observed no increase in English-language activity, either from bots or from human-controlled accounts.

In this edition of *Robotrolling*, we introduce the Global Database of Events Location and Tone (GDELT). This database of news articles helps map how the conversation about NATO in Poland and the Baltics is covered by news media, and serves as a contrast to the environment observed on Twitter and VK. This contrast reveals that in April 2021—as Russian troops mobilized along the Ukrainian border—inauthentic Russian accounts were also disproportionately active online.

We round off the issue with a discussion of how AI can help us better understand the global news environment in near-realtime, based on conversations with StratCom COE expert Gundars Bergmanis-Korāts and GDELT-founder Kalev Leetaru. ■

The Big Picture


In this edition of *Robotrolling*, we continue to track manipulation on social media platforms about the NATO presence in Poland, Estonia, Latvia, and Lithuania. Our analysis focuses on the activities of automated accounts (bots) and coordinated anonymous human accounts. With this—our seventeenth—issue we are making a change to publishing the report biannually, with six-month reporting windows. The period currently under consideration in this spring/summer issue is 1 February to 31 July 2021. This window is compared and contrasted with the previous six months, 1 August 2020 to 31 January 2021.

In the Russian-language information space automated activity has increased markedly in the previous six months. Currently, 36% of tweets come from automated accounts, constituting a 50% increase compared to the previous period. English-language automated activity has also increased, but from a much lower comparative baseline.

Not only was bot messaging more widespread, but we also note a divergence in Russian- and English-language messaging volumes: the number of English messages (excluding retweets) about the NATO presence was unchanged at 11 200; Russian-language messages increased by 40% to 7 200. On VK, the total messaging volume increased from 43 000 to 58 000—a 35% increase. Simultaneously, the percentage of automated Russian-language accounts increased on Twitter from 16% to 19% and on VK from 14% to 16%.

Automated activity on VK this quarter mirrored our findings for Twitter. The proportion of bot activity increased along with message volume. On VK, 56% of messages were posted in groups, an increase from 51% in the previous period.

The increased bot levels for spring and summer 2021 are a statistical outlier. Since starting our observations in 2017 the overall trend has been for gradually declining levels of automated activity on Twitter, thanks to the company's efforts to clean up the platform. The trend has stalled and may now even have reversed. This change in trajectory coincides with the run-up to the Russian military's Zapad exercises to take place in September 2021. As ever, social media companies should not rest on their laurels thinking the battle against bots has been won. ■



Country Overview

While Russian-language bot activity increased in online discussions about all four countries, the largest volume of authentic messaging was provoked by a statement made by President Putin on 9 June. He dismissed the idea that Ukraine might ever join NATO, as missiles placed in Kharkiv could reach Moscow in 7–10 minutes. Social media users were quick to question both the President's inference and grasp of geography, pointing out that NATO-members Estonia and Latvia were closer to St Petersburg and Moscow.

English-language conversations revolved around Belarus' hijacking of RyanAir Flight 4978 on 23 May, the June 2021 NATO Summit in Brussels, and the bilateral meeting between Presidents Biden and Putin shortly thereafter. However, the largest spike in inauthentic activity occurred on 16 June, when numerous accounts questioned the Oxford-AstraZeneca vaccine by highlighting the divergent approaches taken by NATO member states.

Estonia

On 17 May, bots created more than 70% of Russian-language tweets mentioning NATO and Estonia. Several events occurred around this time. English-speaking Twitter users focused on the Italian F-35 jets deployed in support of NATO's Baltic Air Policing mission that scrambled to intercept a Russian military transport plane flying close to the Estonian coast. Russian-language bots highlighted reports of a scuffle involving British soldiers at a fast-food restaurant in the town of Tapa, and promoted the message that the Estonian-led exercise Spring Storm 2021 demonstrated how NATO was amassing troops on Russia's borders. The focus of fake English-language accounts was low and generally not explicitly targeting the NATO presence so much as conversations and hashtags about the Nord Stream 2 project that, when completed, will deliver Russian gas directly to Germany through a pipeline under the Baltic Sea.

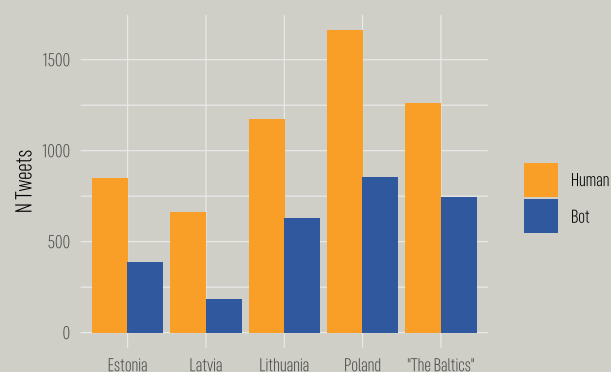


Figure 2: Distribution of Russian-language tweets mentioning NATO and Estonia, Latvia, Lithuania, or Poland.

Latvia

This period Latvia was the subject of the lowest number of messages by both bot and human-controlled accounts in Russian and in English. The single Latvia-specific event to attract any notable bot attention was an interview given in April by Foreign Minister Edgars Rinkēvičs, in which he reiterated Latvian support for NATO Membership Action Plan for Ukraine.

Lithuania

Lithuania, together with Poland, received the most attention in inauthentic Russian-language content this quarter. Also related to Lithuania was the only notable spike present in both the English- and Russian-language environments—the press conference involving Spanish Prime Minister Pedro Sánchez and Lithuanian President Gitanas Nausėda at the air base in Šiauliai that was interrupted as jets scrambled to intercept two Russian military aircraft. Portrayed by some commentators as an embarrassment for the Alliance, others were quick to point out that the incident made a mockery of the Kremlin's narrative of 'NATO the aggressor'. One incident, which notably did not elicit fake activity in either the English or the Russian-language space, was Belarus forcing a Vilnius-bound Ryanair flight to land in Minsk to take opposition journalist Roman Protasevich into custody. President Nausėda called on NATO and EU allies to respond to the threat posed by the Belarus regime.

Poland

In May, English-language bots spread the news that Poland had agreed to purchase Turkish-made drones. The single largest spike in Russian-language activity in this reporting period was primarily caused by bots commenting on an opinion piece written by retired General Skrzypczak. In the article, he discussed NATO's preparedness for information confrontation with Russia and its ability and willingness to defend the Suwalki Gap. Russian bots amplified news stories that absurdly interpreted Skrzypczak's commentary as a blueprint for a NATO invasion of Kaliningrad.




Figure 3: Percentage of bot-generated Twitter messages in English and Russian over the past three reporting periods.

Themes

In The phrase ‘the Baltics’ is a convenient shorthand expression that encapsulates a range of meanings, differing by audience. Bots are no exception: Figure 2 shows that Russian-language posts referencing NATO actions in ‘the Baltics’ are about 25% more likely to be shared by bots than messages naming individual countries. One explanation is that Russian officials regularly paint ominous pictures of NATO actions, unconstrained by a need for accuracy or detail, and that such statements are popular with the bots (or the people who operate them). One prime example, widely amplified in April 2021, was Defence Minister Sergei Shoigu’s claim that NATO was busily relocating troops from the US to Russia’s borders, especially to the ‘Black Sea and Baltic regions’.

The aggregate numbers shown in Figure 3 reveal that the current level of automated activity is a return to the levels seen during the Spring and Summer of 2020. A number of explanations come to mind: first, military exercises, often planned for the spring and summer, have been of recurring interest to Russian news outlets, including those habitually promoted by bots. Second, we have for several years observed a pattern whereby interest in the NATO presence declines during the winter period, especially during the holiday season in December and January. Third, 2021 will see the return of the quadrennial Zapad exercises along the Russian Federation’s western border. The run-up to the previous Zapad exercises in 2017 coincided with the highest level of bot activity we have observed. In April 2021, Russian troops mobilized along the Ukrainian border. It may only be a coincidence that the first significant return of bot activity coincided

with these events, however, it will be interesting to monitor how developments in the information play out in the second half of 2021.

In this issue of *Robotrolling*, we introduce a snapshot of traditional and online media reporting on the NATO presence by using data from the GDELT database. By filtering for articles containing phrases similar to our social media search terms we can create an analogous timeline of related news coverage. This is shown in Figure 4, superimposed over data from Twitter and VK.

Figure 4 shows that in this reporting period the NATO-related events discussed on social media platforms have diverged sharply in the English- and Russian-language environments. There is little correlation between English- and Russian-language sources. For Russian, the main spikes on VK and Twitter are correlated, while for English the spikes on Twitter coincide with activity covered by the English-language news media as captured by GDELT. However, except for spikes coinciding with the detention of Protasevich in May and the interruption of the Sánchez-Nauséda news conference in July, there is little in the graph to suggest a common ground truth for commentary. In particular, the six weeks starting on 1 April show elevated activity levels in both VK and Russian-language Twitter. These levels were equivalent to those observed during the NATO summit and the Biden-Putin meeting in June but at a time when the international news media found little of note to report on the NATO presence. ■




Figure 4: Timeline of VK, Twitter, and GDELT mentions of the NATO presence.

Mapping the information space with GDELT

GDELT (Global Database of Events Location and Time) is an events database. In the words of founder Kalev Leetaru, this ambitious project aims, to help users access “massive firehoses of the world to identify breaking events in realtime, trace the spread of narratives before they become events and understand the flow of ideas, beliefs and narratives around the world, from combatting falsehoods and lending context to helping us better understand each other”.

The database of news-related material is organised into multiple tables. Each table presents the data in a slightly different way, providing a range of use cases. GDELT is based on a stream of digitised news in more than 100 languages, which is standardised using machine learning. Through a collaboration with Google, GDELT applies AI tools to understand the content of the texts it processes. These tools systematically extract key features, such as the places, events, and people mentioned, and also pulls notable quotes, images, and estimates of sentiment.

Figure 5 is a visualisation of the outlets discussing NATO as captured by the GDELT similarity graph. The dataset is new, dating only to 2 July 2021, but the connections already offer insight into patterns within the Russian-language information space. The data consist of pairwise similarity scores for nouns in news texts in 65 languages, standardised through machine translation, making it possible to map similarity patterns within global news output.

The graph shows news outlets as circular nodes. The links between them represent highly similar articles that also mentioned NATO. Mapping the information in this way reveals that many smaller outlets are clustered around larger news agencies. These smaller outlets typically rewrite or republish material in near-verbatim form.

This republishing is significant in at least two ways. First, it shows which sources are most influential. Coverage by news agencies matters because so many publications rely on them for content. What they select gets disseminated widely. Second, patterns of republishing can expose clusters of outlets that act in coordination, amplifying particular messages.

The colouring on the graph is calculated using a community detection algorithm. Very roughly, the orange outlets represent Russian-language sources disproportionately focused on Ukraine. The two blue communities are primarily Russia-focused, with the darker blue tending toward more local news.

It is no surprise to see the news agencies Ria and TASS emerge as important nodes in the graph. More striking is the hub-function performed by outlets such as Riafan (part of the Prigozhin / St Petersburg troll farm media manipulation empire), the conservative and historically anti-Baltic Regnum, and the ‘information agency’ news-front.info.




Figure 5: Russian-language news outlets, connected by similarity scores for articles mentioning NATO (July 2021).

Mapping the information space with GDELT (continued)

GDELT prioritizes scale and speed over precision by design. This is an almost inevitable tradeoff, but continual improvement of web scraping algorithms still offers real potential for refining the system. It is left to the end-user to tidy the data and ensure the results make sense. Sometimes they don't, as the web scrapers may include junk data together with the news. News aggregators that use an infinity scroll are one notable example of such noise. News texts previewed in the sidebar and caught by the scraper are another. Whenever such previews contain the search terms, a new false positive is introduced into the database.

NATO StratCom COE AI expert Gundars Bergmanis-Korāts argues that in practice this tradeoff means that work is required for GDELT to be used as a monitoring tool. In its raw form, noise in the data means timeseries tend to be 'spiky', and subtle variations may have little discernible meaning. Currently GDELT works better as a discovery tool that analysts can use to confirm novel trends. A particularly powerful example of this, given by Leetaru, is how "BlueDot sent out one of the very first worldwide alerts of Covid-19 on December 31, 2019 using their machine learning algorithms to scan GDELT's data streams for disease outbreaks. No-one was watching local Chinese language news coverage in Wuhan in December 2019 looking for a pandemic, but because we machine translate everything we see, we were able to pick up those earliest first glimmers of Covid-19."

A challenge working with GDELT and other AI-tools is that they are developed primarily for English-language content. There is a lot of hope for language agnostic AI models, but at the moment Leetaru sees machine translation and English-language models as preferable; there is still a large gap between the promise of language-independent systems and reality.

To Bergmanis-Korāts, GDELT is a powerful project with great future potential. It continues to improve technically and is updated in near-realtime, enabling access to scraped and post-processed data without worrying about using NLP or computer vision algorithms. Forget about building custom web scrapers. In addition, post-processing and feature extraction tasks such as machine translation, named entity recognition, and sentiment analysis are already done. Such data appear invaluable, but bear in mind that identifying the correct dataset slice requires some specialist skills such as understanding the datasets and features they contain, building proper SQL queries using Google's BigQuery interface, and performing data cleaning and further processing/visualization. GDELT is a powerful platform for either monitoring or discovering new previously unseen relations and patterns in the multi-lingual digital news space. ■

Author: Rolf Fredheim

With thanks to Kalev Leetaru and Gundars Bergmanis-Korāts.

**NATO STRATEGIC COMMUNICATIONS
CENTRE OF EXCELLENCE**

The NATO StratCom Centre of Excellence, based in Latvia, is a multinational, cross-sector organization which provides comprehensive analyses, advice and practical support to the Alliance and Allied Nations.