

ROBOTROLLING
2019. ISSUE 2



ROBOTROLLING

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



Executive Summary

On the popular Russian-language social network VK, material about the NATO presence in the Baltics and Poland was viewed no less than 11 million times this quarter (February – April 2019). 93% of these views were for material from community spaces.

On VK, community spaces are increasingly important, both as a sources of content and as places for discussion. The move to groups has implications beyond the Russian-language space. Facebook has recently launched a push to promote community spaces. These spaces, normally closed to researchers, offer huge potential for misuse and manipulation.

Our investigation of VK community spaces reveals that the vast majority of groups in which NATO is discussed are communities with radical pro-Kremlin or nationalist tendencies, or dedicated to the conflict in Ukraine. These communities generate more

posts and attract more views even than communities created by Russian state media outlets.

On Twitter, bots tweeting in Russian remain a bigger problem than bots tweeting in English. In Russian, they account for 43% of all messages—a significant increase in recent months. In English bots posted 17% of messages. English-language bots this quarter overwhelmingly amplified news content from RT (formerly Russia Today) and other pro-Kremlin news outlets. On all platforms, discussion regarding NATO troops in Poland attracted the largest number of posts this quarter.

Finally, in this issue we publish our first case study of manipulation on Facebook. It looks at the degree to which bots and trolls targeted posts promoted by Latvian political parties contesting the European Elections in late May 2019.

The Big Picture

This issue of the quarterly report Robotrolling continues to track trends in the manipulation of information regarding the NATO presence in the Baltic States and Poland on the social media platforms Twitter and VK. Manipulation takes two main forms: through automated accounts (bots) and coordinated, anonymous human accounts.

During the period February 1 to April 30, we observed an increase in the volume of social media messaging mentioning the NATO presence in the Baltics and Poland, as well as in the number of users engaged in the conversation.

The number of posts on VK was 26 000, a 20% increase over the previous quarter. For Twitter the number of posts was 9 300, of which 6 300 were in English. This constitutes a 25% increase for English messaging and a 45% increase for Russian.


The number of Twitter-users posting in Russian about the NATO presence in the Baltics and Poland has increased substantially this quarter, whereas the number of active English users is 6 % lower.

In Russian-language Twitter we estimate that activity from automated accounts (bots) also increased by seven percentage points to 43% of all messages.¹ Our algorithm identifies 10% of Russian-language users as normal human accounts. This figure is comparable to the previous quarters. Consequently the space remains dominated by bots, anonymous users, news aggregators, and various institutional accounts.

¹ At first glance, the bot percentage of Russian-language Twitter users appears to have decreased. However, this is only due to our algorithm reclassifying a number of accounts previously identified as bots to anonymous, human-controlled accounts.

Automated activity on English-language Twitter has remained at a stable level compared to previous quarters. The same is true for VK (Russian only), where the percentage of bot users is unchanged at 17%. These bots posted 35% of messages on VK, not including messages in community spaces.

Groups account for more than 60% of all posts, an increase of 5% over last quarter. Moreover, 93% of all views on VK were for material from community spaces. The total number of views on VK for material about the NATO presence was no less than 11.2 million. ■



Country Overview

Bots remain a consistent source for Russian-language messaging regarding the NATO presence in the Baltics and Poland—both on Twitter and VK.

VK and Twitter activity was high at the beginning of April, coinciding with the 70th anniversary of NATO. Over 60% of the messaging on VK about the anniversary took place in groups, whereas on Twitter the discussion was dominated by anonymous users focusing on events in Poland and the Baltic region in general, paying least attention to Latvia.

Overall, human users on VK posted similar levels of messages to bots, both prioritising Poland. VK human users' activity tended to be about Estonia and Lithuania, whilst automated posts were more focused on Estonia and Poland. Likewise, VK group activity was also higher for Poland.

Automated accounts actively amplified content from RT. They did so directly—by sharing links to the Kremlin-owned media outlet—and indirectly—promoting content from Kremlin-aligned and anti-Western media that reposted RT articles (e.g. *Voice of Europe*).

Estonia

Estonia received the second highest number of mentions and was the main target of activity from anonymous accounts on Twitter. The Spring Storm military exercise attracted higher levels of interest on 29 April, especially in Russian-language Twitter where bots posted 51% of all the messages. VK groups were particularly active, creating around 70% of the posts about the exercise. Comments about the NATO military presence in the Baltics on VK peaked on 15 April, with humans and groups producing more than 70% of the messages.

Latvia

Latvia is the only country where the volume of Russian-language posts on Twitter was higher than the volume of English-language posts. The conversation was on average more muted than in the other countries, even at the time of NATO's 70th anniversary. However, comments about the Iron Spear military exercises conducted in Ādaži attracted high levels of Russian-language bot activity on Twitter on 8 April.

Lithuania

Russian-language Twitter and VK activity peaked on 18 April after allegations of NATO frigates entering the Baltic Sea. This event produced three times more activity than any other during the studied period. Here automated accounts on VK were unusually active, creating 37% of the volume of conversation (excluding groups). In general, anonymous accounts were disproportionately active, generating a whopping 58% of all Twitter posts.

Poland

This quarter, Poland received the highest number of mentions across VK, Russian-, and English-language Twitter due to three key events that raised the level of attention mainly from anonymous accounts on Twitter and VK groups: US Secretary of State Mike Pompeo's visit to Poland on 13 February; the 20th anniversary of Poland's accession to NATO on 12 March; and the reports on 23 March of NATO's support in building a US military storage facility in Poland.

VK group messaging about Poland remains high, constituting 65% of the total mentions of NATO and Poland on VK. On VK, bots and anonymous users posted more about Poland than the three Baltic States put together. ■




Figure 2: Total number of posts from automated (bot) and anonymous accounts, divided by country. For VK, only posts to users' timelines (TL) have been included.




Figure 3: Total view counts for posts in the most popular VK groups and pages, divided by orientation as identified in the Robo-topics section

Themes

Our analysis reveals that VK community spaces—groups, pages, and communities—are remarkably impactful. This quarter, content shared on VK about the NATO presence in the three Baltic States and Poland attracted at least 11.3 million views.² 93% of these views were for content shared in community spaces.

Three community spaces embedded within the Russian nationalist and pro-Donbass community spread content which taken together was viewed more than 2.5 million times. These groups were world_policy, strelkov_info, and antimaidan. The most popular groups can be seen listed in Figure 4.

This quarter English-language content from Kremlin-backed RT, was heavily amplified by bots and anonymous users promoting the narrative of NATO as an aggressive power committed to threatening Russia. Coverage of NATO-related military activity near the Russian border was persistent. Poland and the Baltic nations were depicted as enabling NATO’s intimidating presence in the region. Platforms that republished RT content, such as the English-language outlets *Voice of Europe* and *Infowars*, had their articles promoted by bots, thus extending the reach of Kremlin-controlled media.

We note some differences between English-language and Russian-language bots: The former were committed primarily to the amplification of RT content; the latter drew on articles from different Russian-language sources, facilitating the dissemination of the narrative of ‘NATO as an aggressive power that builds tensions with Russia and risks destabilising the global order’. This activity appears to fit into three overarching objectives: the promotion of Russian victimisation, denouncing NATO’s Russophobic strategies, and undermining the cohesion of Western alliances.

Lately the NATO presence in Poland has been of particular interest, both to key VK groups, and to bots and automated accounts on Twitter. This quarter, social media activity targeting Poland increased across all platforms and languages. English-language messaging on Twitter accounts peaked after the visit of US Secretary of State Mike Pompeo to the multinational battlegroup in Poland as part of NATO’s Enhanced Forward Presence (eFP) and with the 20th anniversary of Poland’s accession to NATO. By contrast, Russian-language messaging on Twitter peaked on 23 March, amid reports of a plan to build a \$230 million storage facility in Poland.

In mid February, following Secretary Pompeo’s visit to Poland, accounts shared an article from RT which cited the Polish Foreign Minister’s concerns about the creation of a European Army in isolation from NATO and the presence of American troops in Europe. The article questioned the credibility of EU efforts to foster greater internal military cooperation. However, it was also disseminated by accounts to express Russia’s concern regarding increasing tensions due to increasing US-Polish military collaboration. Whilst some bot accounts simply amplified content to increase visibility, others combined this with anti-US, anti-Western messages, for example: ‘Poland always con-mitts [sic] themselves to the wrong partner on the wrong moment.’

On 8 April, accounts and websites republishing material from *Sputnik* reported on Iron Spear 2019. However, a typographical error, probably originating in a Russian-language news wire report, was reproduced by dozens of news outlets and many more bots, resulting in a spike of references to a military exercise with the improbable name ‘Iron Spread’. Indiscriminate amplification of pro-Kremlin content is not without its perils; fact-checking evidently is not a priority. ■

² View counts given at data collection time. Collection occurs daily, meaning the final view counts will be higher.




Figure 4: Timeline showing volume of posts on Twitter and VK.

Robo-topics

For our analysis of groups on VK we identified a core set of 4000 communities posting about NATO. To model the space, we calculated which communities had the most similar memberships. From these results we created a network visualisation in which communities' positions were determined by their similarity to each other. Next we used a clustering algorithm to group the communities. We manually inspected the resulting clusters and labelled them according to their overall orientation. In cases where group clusters were similar we further combined groups. For instance, the clusters of groups thematically relating to Armenia, Georgia, Belarus, and Murmansk were all combined into the category 'International', which itself was later included within 'Other news'.

Figure 5 below maps the resulting 'groups of groups'. In it, group clusters are presented as nodes, whereas edges represent communities with similar memberships to other clusters. Thus the large collection of groups exhibiting predominantly nationalist leanings is closely related to the many communities with a focus on the so-called Novorossiia region, that is, the self-proclaimed republics of Donetsk and Luhansk.

In the figure, the size of clusters is scaled to reflect importance according to three different metrics: total number of groups included within a cluster, the number of posts mentioning NATO, and the total number of views these posts received.

Here the 'other news' category, which includes many groups aggregating news, often from a regional perspective, contains the largest number of groups. However, the number of posts is in the middle of the range, and the view count shows that posts in these communities rarely attract a big readership. Conversely, the number of groups and posts within the 'State media' cluster is very small. It contains groups based around media outlets such as RT and *Ria Novosti*. However, this content received millions of views and has much broader significance and impact than material from the other clusters. The small number of oppositional or independent groups, centered around news outlets such as Meduza was disproportionately influential.

The vast majority of posts, however, originate from groups within the clusters labelled Novorossiia and Nationalist++. The ++ in this label indicates that the community is largely nationalist, but that some other sub categories such as orthodox and pro-Kremlin have been included. Not only is there a large number of communities spreading huge volumes of information, they are also influential with view counts in the millions. 50% of all views of material about NATO shared on VK took place within communities with radical pro-Kremlin or nationalist tendencies, dedicated to Novorossiia, or to weaponry and other military matters. State media attracted a further 26% of views.

Remarkably, in this environment, the space occupied by pro-Kremlin media is comparatively moderate. ■

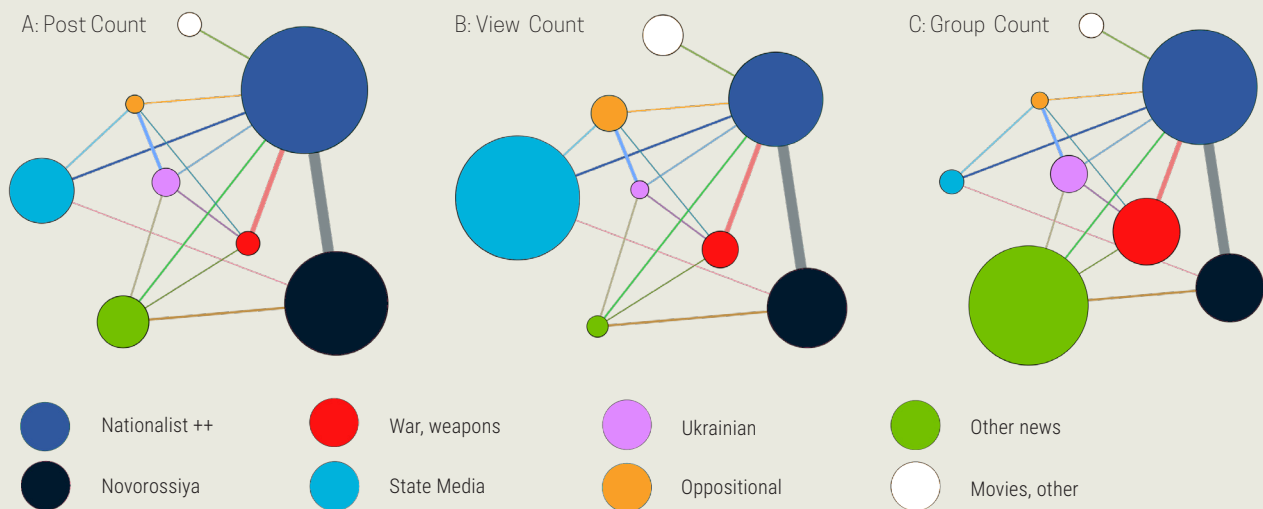


Figure 5: Network map showing groups of community spaces, scaled according to: the number of posts mentioning NATO, the total number of views for these posts, and the number of communities included in a cluster.

Case Study: Trolling ahead of European Parliament

Prior to the European Parliamentary elections, held 23–26 May, the NATO Strategic Communications Centre was approached by representatives of a Latvian political party expressing concerns about possible interference targeting their social media campaigns on Facebook. In response we conducted a quick investigation of activity contesting the European elections around posts promoted by the six most popular parties in Latvia.

Four of the six most popular political parties in Latvia ran Facebook advertising campaigns ahead of the European Parliament elections. Campaigning techniques varied, with some parties boosting only their Facebook pages and others boosting their top candidate pages. 63% of EU election ads were run from the candidates' own pages.

We employed the same account typology when studying Facebook as for VK and Twitter analysis, dividing profiles into recognisably human, anonymous human-controlled, automated, institutional, and news. Our investigation was conducted manually, due to restrictions imposed by Facebook on researchers using automated data collection.

Our analysis did not reveal any evidence of foreign interference. It showed no obvious use of automated accounts to comment on posts. We did not explicitly investigate whether bots were used to inflate likes, but based on the information available, we consider this did not happen on a large scale.

When campaigning, parties frequently use Facebook's advertising platform to promote posts and reach a larger audience. Unless audiences are carefully targeted, such posts are liable to attract

hostile responses from citizens who support other political parties. Content on all parties' pages attracted negative, inflammatory, or off-topic comments. Moderators frequently hide offensive comments. The volume of abusive or off-topic comments, though significant, was not at such a scale that it overwhelmed moderators or clearly skewed engagement metrics.

In total, 42% of all comments on promoted posts were hidden by page administrators. All parties hid roughly one third of the comments. While this figure is large, it need not point to an organised campaign. A portion of the hidden content consisted of short, offensive comments written in simple Russian. However, the accounts that posted these comments appear authentic and their activity was not coordinated.

In total, 15% of comments (that survived moderation efforts) came from anonymous and likely inauthentic accounts. Some topics attracted greater levels of abuse: candidates' past achievements and future aims, the Kremlin's information warfare, and the corruption case against the former mayor of Riga. No single party was targeted more than the others.

A handful of likely inauthentic accounts systematically commented on promoted posts of all parties. These accounts do not seem to be connected or coordinated, as their profile appearance and engagement tactics differ. The accounts which commented most intensively on party posts may have sought to suppress voter turnout; their posting followed one unifying theme, namely that voting in the EP elections was pointless and that Latvia has no influence in the EP. ■

Prepared by Dr. Rolf Fredheim, Nora Biteniece and Belén Carrasco Rodríguez published by

**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.

www.stratcomcoe.org | [@stratcomcoe](https://twitter.com/stratcomcoe) | info@stratcomcoe.org