Article

new media & society
1–22

DOI: 10.1177/14614448241254126
journals.sagepub.com/home/nms
Sage

Protest and repression on social media: Pro-Navalny and pro-government mobilization dynamics and coordination patterns on Russian Twitter

Aytalina Kulichkina
University of Vienna, Austria

Nicola Righetti
University of Urbino Carlo Bo, Italy

Annie Waldherr
University of Vienna, Austria

Abstract
In this study, we examine connective protest mobilization and suppression during the 2021 protests in Russia. We use time series analysis to study the dynamic interplay between the pro-Navalny movement and pro-government countermovement on Twitter, complemented by network analyses of co-retweeting networks to assess the movements' coordination patterns. Findings show that pro-Navalny accounts were more active and coordinated within more centralized Twitter networks than pro-government accounts. Contrarily, the pro-government camp employed preventive communication tactics and coordinated in more clustered networks. Granger causality tests reveal that pro-Navalny tweeting activity triggered increased pro-regime reaction during the largest protests on 23 January and 21 April, whereas pro-government tweeting activity caused the escalation of pro-Navalny reaction during the 14 February protests. Both sides' tweeting activity decreased after the February protests, presumably due to external repression. These findings contribute to a deeper understanding of online mobilization and coordination strategies via social media in authoritarian contexts.

Corresponding author:
Aytalina Kulichkina, Department of Communication, University of Vienna, Kolingasse 16, 1090 Vienna, Austria.
Email: aytalina.kulichkina@univie.ac.at

Keywords
Connective action, coordination, networks, protest, repression, Russia, social media, Twitter

Introduction

The scholarly discussion on social media's role in political protests within authoritarian regimes unfolds as a dynamic dialogue between two distinctive perspectives. One underscores the empowering capacity of social media for protest mobilization, including information sharing, coordination during demonstrations, or internationalization of protests (e.g. Bruns et al., 2013; Lotan et al., 2011; Tufekci and Wilson, 2012; Tucker et al., 2016). The other focuses on the suppressive uses of social media or digital repression techniques, such as content filtering, online propaganda, disinformation, or digital surveillance (e.g. Earl et al., 2022; Feldstein, 2021; Hellmeier, 2016; Morozov, 2011; Xu, 2021). Taken together, they point to the anticipated "cat-and-mouse dynamics" between citizens and governments due to technological change (Kendall-Taylor et al., 2020: 115), highlighting the need for empirical investigation of this phenomenon.

Our study contributes to both strands of literature by employing Bennett and Segerberg's (2013) connective action framework to conceptualize and examine protest-related action and counteraction in an authoritarian context. Specifically, we aim to integrate social media's potential to both empower and constrain protests that can influence the development of a protest movement in a repressive regime. Moreover, we investigate the role of coordination in connective action, discerning three related yet distinct characteristics of coordinated networks: synchronization, centralization, and modularity. These characteristics provide valuable insights into different mobilization strategies and networked organization logics within social media platforms. In authoritarian settings, where participation in protest movements is associated with risks, understanding the scope of coordination for both protest and repression purposes is crucial. This aspect has not received sufficient scholarly attention thus far, and our study aims to push forward the discussion by employing a case study approach.

We examine the dynamics of the 2021 Russian protests on Twitter, with particular attention to the interplay between online activism and offline events, as well as the coordination efforts in the pro-Navalny and pro-government camps on the platform. To this end, we employ two distinct methodological approaches. First, we conduct a time series analysis to investigate how the tweeting patterns of pro-Navalny and pro-government accounts relate to protest events and each other. Second, we employ network analysis to understand the characteristics of coordinated networks of both camps on Twitter. Our study aims to comprehensively examine these dynamics within a single framework, which we believe has not been attempted before. Moreover, our methodological contribution includes the development of a tool that can identify coordinated networks based on digital traces. Through our analysis, we hope to provide valuable insights into evaluating counter-dynamics of digital protest and repression, ultimately contributing to a more nuanced understanding of connective action in authoritarian contexts.

Connective action and connective counteraction

Research on information communication technologies and protest has identified various categorizations of protest organization methods enabled by digital media. This study relies on the widely recognized analytical framework of connective action proposed by Bennett and Segerberg (2013), which we find applicable to the complex nature of protests in authoritarian settings. The framework encompasses three forms of action that can be viewed as broad analytical categories: (1) organizationally brokered collective action, (2) organizationally enabled connective action, and (3) crowd-enabled connective action, where the former follows a more conventional protest organization logic and the latter two utilize digital media as an essential tool for network-building and participation (Bennett and Segerberg, 2013).

Specifically, organizationally brokered collective action is characterized by strong organizational coordination, social technology use by organizations for coordination, collective action frames, organizational management of social networks, and organizations in the foreground (Bennett and Segerberg, 2013). On the other hand, organizationally enabled connective action features loose organizational coordination of action, provision of social technology by organizations, organizationally generated inclusive personal action frames, organizational moderation of personal expression, and organizations in the background in loosely linked networks. Finally, crowd-enabled connective action is distinguished by little or no formal organizational coordination, large-scale personal access to social technologies, inclusive personal action frames, sharing personal expression over social networks, and shunning the involvement of organizations by crowd networks.

This analytical lens has been applied to study social media-enabled network formation and organization of various protest movements worldwide. Relevant to our case, Toepfl (2018) scrutinized previous large-scale Russian protests of 2011–2013 and found that the protests initially utilized organizationally enabled connective action, with activists heavily leveraging digital media. However, the protests eventually transitioned into more conventional organizationally brokered networks where digital technologies played a less significant role. This insight is crucial since it emphasizes the dynamic nature of protest movements and the framework's applicability to different movement stages. Nevertheless, the connective action framework has not yet been utilized to simultaneously examine counteractive networks aimed at suppressing protests, especially in authoritarian settings.

Authoritarian governments have adapted to using digital tools to prevent protests, suppress dissent, and endorse regime supporters (e.g. Feldstein, 2021; Weidmann and Rød, 2019). Such efforts require the involvement of different actors, ranging from state agents tightly connected with political officials to private agents operating autonomously (Earl et al., 2022). Ordinary citizens may also voluntarily engage in digital campaigns supporting the status quo. Research indicates that individuals, especially those from disadvantaged groups, often justify and defend their political system, even when it goes against their interests (Jost et al., 2004). Therefore, the emergence of connective counteraction networks in response to protesters' connective action is a logical and anticipated phenomenon in authoritarian settings.

Counteractions may take the form of "political astroturfing—a campaign in which participants appear to be part of a genuine grassroots movement or sentiment, while it is in fact orchestrated centrally and top down" (Keller et al., 2020: 1). They can also involve ordinary citizens encouraged by pro-government organizations, such as political parties, government-owned institutions, or companies. In addition, genuine bottom-up grassroots movements may emerge from individuals who genuinely support the government and justify the regime. Thus, Bennett and Segerberg's (2013) connective action framework can be employed to analyze such counteraction networks, whether they are organizationally brokered by the government or related organizations, organizationally enabled when the government or related organizations only endorse existing supporters, or self-organized by dispersed regime devotees.

It is worth noting that while our study applies this framework to analyze pro-regime and pro-opposition networks in an authoritarian setting, connective counteraction can also occur in polarized democracies. For example, a protest movement advocating for the expansion of rights for historically disadvantaged groups may prompt a countermovement aiming to preserve the status quo and suppress opposition. Conversely, a movement against social change may encounter a countermovement advocating for it. The distinctive aspect of an authoritarian context lies in restricted political freedoms, which limit protesters' abilities to mobilize widespread support and coordinate actions.

Characteristics of coordinated networks

In the connective action framework, all three categories of network organizations involve different degrees of coordination, as discussed earlier. Coordination has been studied as a beneficial strategy for organizing and empowering political protests through various affordances of social media (e.g. Barberá et al., 2015; Bennett and Segerberg, 2013; Bruns et al., 2013; Earl and Kimport, 2011). Notably, using specific hashtags for coordination has been found to increase the scale of protests the next day (Steinert-Threlkeld et al., 2015). More recently, coordination has also been studied as a malicious strategy used for manipulating opinion and spreading disinformation (e.g. Giglietto et al., 2020b; Howard et al., 2018; Marwick and Lewis, 2017). In authoritarian contexts, it can also be employed strategically to divert attention from collective action and the accompanying grievances (King et al., 2017).

While recent discussions highlight malicious motivations behind coordinated behavior on social media, we argue that coordination can serve both connective action and counteraction, encompassing both mobilizing and repressive purposes. Therefore, we adopt a neutral definition of coordination: "the act of making people and/or things involved in organized cooperation" (Giglietto et al., 2020b: 6). Within the connective action framework, coordination degree reflects the involvement of organizational actors in the underlying networks, making the structural characteristics of these networks a defining element of the type and even outcome of connective or collective action (Bennett and Segerberg, 2013).

In this study, we propose focusing on three related but distinct characteristics of coordinated networks: (1) *synchronization*, (2) *centralization*, and (3) *modularity*. Synchronization refers to the proportion of social media accounts involved in coordinated action over a period of time. It assesses the extent to which social media posts result from organized coordination and the efficiency of coordinated action. High levels of synchronization within short time windows suggest intensive coordination and substantial organizational involvement to carry it out. Centralization measures the extent to which a coordinated network is dominated by a single actor based on its position in the overall network. High centralization scores reflect the presence of highly involved leaders orchestrating coordination. Modularity is another feature that assesses the degree to which a coordinated network is structured into smaller subnetworks or communities. High levels of modularity within short time windows may suggest the involvement of multiple teams participating together in coordinated behavior.

Together, these characteristics can provide insights into the nature of connective or collective action. For example, a crowd-enabled network would have low levels of synchronization, centralization, and modularity, as grassroots mobilization is decentralized and "travels easily across large and diverse populations" (Bennett and Segerberg, 2013: 37). Contrarily, high levels of synchronization, centralization, and modularity would indicate a more conventional collective action, since they rely on strong organizational coordination and place greater emphasis on interpersonal networks (Bennett and Segerberg, 2013). Finally, moderate centralization and modularity would suggest an organizationally enabled network structure with potentially varying degrees of synchronization.

Pro-Navalny and pro-government mobilization on Twitter

We apply these measures to Russia's 2021 pro-Navalny protest movement, a significant milestone in the struggle for political change within a highly affluent repressive regime. On 23 January 2021, tens of thousands of protesters gathered in Moscow, St Petersburg, and across Russia to rally in support of the arrested opposition activist Alexei Navalny. The preparations for these large-scale demonstrations began just a few days before when Navalny and his team called on their supporters to take to the streets via social media. Although the government warned against attending the rallies, the call spread nationwide and was endorsed by thousands online. However, pro-government actors and organizations used the same social media platforms to discourage the public from attending the rallies. Thousands of trolls and bots targeted opposition figures and independent Russian media online to hinder their work and contain the protests (Baklanov, 2021; Sobol, 2021). Despite these obstacles, the demonstrations repeated on 31 January and 2 February, continued in courtyards on 14 February, and culminated in nationwide street protests on 21 April.

The protests unfolded during the COVID-19 pandemic, which saw increased use of digital media due to physical social distancing and self-isolation measures. This made it inevitable for protesters to use social media for mobilization while the regime attempted to suppress dissent online. After the 2011 large-scale protests enabled by social media, Russian authorities intensified efforts to regulate the Internet, which resulted in increased risks associated with expressing anti-regime views on social media (Lonkila et al., 2021). Although various social networking sites were still available in 2021, the choice of platform for mobilization was a safety concern for anti-regime activists. Navalny and his team utilized various platforms where they had previously gained a substantial following, including globally recognized Twitter, Facebook, Instagram, YouTube, Telegram, and TikTok, as well as Russian-based VKontakte and Odnoklassniki (Glazunova, 2022). However, Russian-based platforms pose higher risks for the broader public than foreign social media that do not store user data in Russia. Foreign platforms offer lower-risk infrastructure and a favorable environment for leadership-critical publics, allowing users to explicitly voice their criticism toward the country's political leadership (Toepfl, 2020). In contrast, local social media services such as VKontakte or Odnoklassniki are known for sharing personal user data with the government (Poupin, 2021), deterring leadership critics and attracting regime supporters or politically disengaged individuals.

While global messengers like Telegram have gained increasing importance for mobilization in Russia, Twitter stood out as a significant concern for the government, as evidenced by its deliberate slowdown by authorities in March 2021 (Xue et al., 2021). Despite its relatively modest user base in Russia, totaling 4.65 million users in 2021 (Statista, 2023), Twitter serves as an important platform for both pro-government and oppositional groups involved in political discourse (Spaiser et al., 2017). Notably, at the time of Navalny's detention, his Twitter account had 2.41 million followers, while his team's account had over 130,000 (Navalny, 2021; Team Navalny, 2021). Considering that digital repression tactics within this leadership-critical environment may differ from those employed on other platforms, Twitter is particularly noteworthy for examination. Specifically, it offers a unique opportunity to explore the dynamic interplay between online protest and repression, especially given that the public data could be obtained in full via the Twitter application programming interface (API) for academic research. Thus, this article focuses on pro-Navalny and pro-government actions observable on Twitter.

Based on the discussion above, we raise the following research questions:

> *RQ1*: How do the tweeting patterns of pro-Navalny and pro-government accounts relate to protest events and each other?

> *RQ2*: What are the synchronization, centralization, and modularity characteristics of pro-Navalny and pro-government coordinated networks on Twitter?

Method

Data collection

We collected Twitter data through Twitter API v2 for Academic Research,[1] which provided access to the Twitter full-archive endpoint (Tornes, 2021) starting in April 2021. The initial data included Russian-language tweets with protest-related hashtags published from 1 January to 18 December 2021 (when we stopped data collection due to the fading of protest-related tweets). To identify relevant hashtags, we manually monitored Twitter trending topics in Russia daily, starting on 17 January 2021 (the day when Navalny was arrested). To this end, we used the list of trending topics provided by

Twitter's website, occasionally supplemented by references to News and Trending Topics in Russia (2021). Russian-language hashtags connected to the protest movement were then used as search queries to collect data from Twitter. The rationale behind this approach was to obtain posts meant to be seen by users based in Russia.

The hashtags were arranged into two groups, pro-Navalny or pro-government, depending on their framing and the messages tweeted alongside them. Once we had collected data using these hashtags, we identified the 300 most retweeted tweets from both data sets and extracted additional protest-related hashtags from them for further data collection. Overall, we collected 729,246 pro-Navalny tweets and 41,642 pro-government tweets using all identified hashtags (see Supplemental Appendix A for details).

We further systematically validated random samples of 500 tweets from each pro-Navalny and pro-government data set by assessing their relevance to protests and alignment with the stance expressed via the hashtags. The results suggest that the proportions of potentially irrelevant or contradictory tweets are below 5% in both data sets, and the average number of co-retweets they gained is insignificant, as tested with Welch's *T*-tests (see Supplemental Appendix B).

Data analysis

To address RQ1, we examined the time series of pro-Navalny and pro-government Twitter communication throughout the protest movement.[2] We analyzed peaks and troughs with regard to the offline protest events and government actions reported by verified news media. In addition, we provided examples of the most retweeted tweets published during the analyzed periods. Subsequently, we segmented the time series into substantive phases of the protest movement using a data-driven approach, namely, changepoint analysis (Killick and Eckley, 2014: 2; see Supplemental Appendix C).

For each of the nine identified phases, we examined the relationships between the pro-Navalny and pro-government series by fitting vector autoregression (VAR) models followed by the Granger (1969) causality test (see Supplemental Appendix D for details). As stationarity is a prerequisite for the Granger causality test, we applied the differencing procedure whenever necessary based on the augmented Dickey–Fuller (ADF) test results, which is a common approach for achieving stationarity in time series analysis (Shin, 2017: 33; see Supplemental Appendix E).[3]

Next, we fitted the VAR models to each protest phase and performed standard model diagnostics. VAR lag order information criteria yielded different results across phases, ranging from 1 lag for the 14 February protests to 16 for the phase right after. We opted for a lag length of 16 (corresponding to 4 hours) for all phases for clarity of interpretation and because it produced the lowest levels of residual serial autocorrelation. A similar approach was adopted by Freelon et al. (2018) and Bastos et al. (2015), who argue that an identical lag order provides better interpretability and comparability.

We examined the structural stability of the models with the ordinary least square-cumulative sum (OLS-CUSUM) method, heteroscedasticity with the autoregressive conditional heteroscedasticity-Lagrange multiplier (ARCH-LM) tests, and multivariate normality with the Jarque–Bera tests (see Supplemental Appendix F).[4] While all models were structurally stable, many violated heteroscedasticity and non-normality assumptions. To address this, we used a wild bootstrapping procedure for Granger causality tests (Hafner and Herwartz, 2009). In addition, to validate the robustness of our results and approach, we conducted the analysis using the optimal lag length for each series determined by information criteria and diagnostic tests. We also performed Granger tests following the standard procedure instead of the bootstrapping method. This approach produced results largely consistent with our original findings reported below.

Since Granger causality analysis only identifies the presence or absence of a causal relationship between variables, we performed an Impulse Response Function (IRF) analysis (Lütkepohl, 2006) to obtain information about the direction and strength of the identified causal relationships. While Granger causality analysis suggests that one variable Granger-causes the other, IRF analysis shows the extent to which a shock to the first variable impacts the second variable over time. Hence, we combined both methods to understand the relationships between the series better.

To address RQ2, we developed an R application aimed at identifying coordinated accounts. Specifically, we based the analysis on "co-retweeting" patterns within short time frames, following the approach by Keller et al. (2020). Drawing on the literature on coordination detection (Giglietto et al., 2020a; Graham and QUT Digital Observatory, 2020), we performed a network analysis where users were represented as nodes, with links established between nodes when users retweeted the same tweet a certain number of times within a predefined time threshold. We analyzed co-retweeting activities across 15 intervals ranging from 1 second to 1 hour, including accounts that co-retweeted the same tweet from at least once to at least ten times. Specifically, we assessed three coordination measures of the networks: synchronization, centralization, and modularity. Synchronization was measured by calculating the ratio of the coordinated accounts to the total number of accounts, while centralization and modularity were calculated using the general method for calculating network-level centralization scores based on node degrees (Wasserman and Faust, 1994) and the Louvain method for modularity scores (Blondel et al., 2008).[5]

Results

Pro-Navalny and pro-government mobilization dynamics

RQ1 asked how Twitter activities of pro-Navalny and pro-regime accounts relate to protest events and each other. A comparison of tweet volume alone revealed that pro-Navalny accounts were 17.5 times more active than their pro-government counterparts. Figure 1 illustrates the difference in tweeting patterns between pro-Navalny and pro-regime accounts throughout the protest movement. Higher attention peaks on both graphs align with protest days or the days immediately before and after. It is important to note that Russia spans 11 time zones—Twitter users in Kamchatka may be awake and tweeting on protest days while it is not even midnight in Kaliningrad. Nonetheless, the graph indicates that despite increased opposition activity close to and during protest days, the pro-government camp maintained high tweeting activity between the protests on 23 January and 31 January. A similar tendency is observed before the courtyard protests on

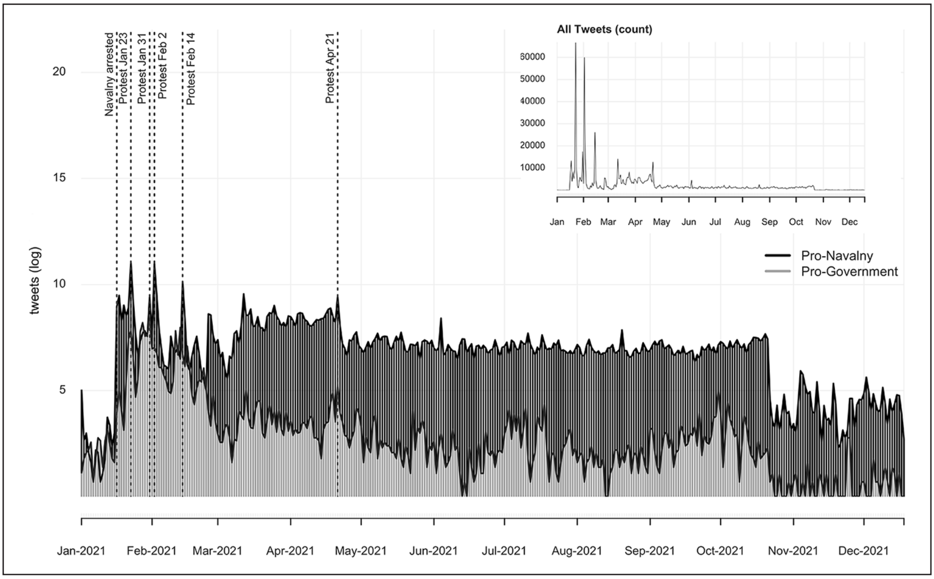

Figure 1. Time series of sampled tweets per camp per day (1 January–18 December).
The series show pro-Navalny (dark gray, $N = 729{,}246$) and pro-government (light gray, $N = 41{,}642$) tweet counts on the log-transformed *y*-axis. Peaks are annotated with relevant events. A plot in the upper-right corner shows the temporal distribution of tweets in absolute numbers.

14 February, with a surge in preemptive pro-regime tweets occurring during a period of reduced activity from pro-Navalny accounts.

A possible explanation of this tendency could be the regime's concerns about upcoming protest waves and their effort to preemptively suppress them by outvoicing protesters on Twitter. The most retweeted posts by the pro-government accounts in the hours leading up to the protest days mainly consist of encouraging messages and videos for the riot police. Although ostensibly intended for law enforcement officers, they could also serve to intimidate protesters. In contrast, the most widely shared pro-Navalny posts from the same time provide information about the upcoming protests and offer advice and warnings on behavior during potential detainment. Notably, the latter could be provoked by the pro-government narrative and could similarly discourage protest participation.

Compared to the protests in January and February, there was notably less Twitter activity in both camps leading up to, during, and after the protests on 21 April, despite their large nationwide scale. This decline could be attributed to the increasing number of protest-related arrests, limited access of protesters to Twitter, and a reluctance to take the risk of posting protest-related content. The most popular tweets before and during 21 April were primarily posted by Navalny's team and friends, calling for urgent medical assistance for Navalny and his immediate release. The pro-government tweets, on the other hand, once again expressed support for the riot police and threatened the public with detentions. Nevertheless, the graph shows that pro-Navalny accounts were remarkably more active overall than pro-government accounts throughout the entire protest movement.

Table 1. Granger causality tests.

Phase	Nav → Gov (*F*)	Nav → Gov (*p*)	Gov → Nav (*F*)	Gov → Nav (*p*)	Instant. (chi-sq.)	Instant. (*p*)
Navalny arrested	1.275	.342	0.937	.460	1.212	.271
Protest 23 January	1.873	**.050**	1.708	.098	2.960	.085
Between protests 1	1.020	.387	1.019	.493	6.135	**.013**
Protest 31 January	1.284	.159	1.050	.167	0.476	.490
Protest 2 February	1.075	.398	0.495	.838	0.124	.725
Between protests 2	1.328	.178	1.456	.128	3.868	**.049**
Protest 14 February	0.193	.407	2.513	**.045**	1.418	.234
Between protests 3	1.717	.086	0.995	.546	0.005	.944
Protest 21 April	2.212	**.001**	1.232	.126	5.379	**.020**

Cell entries are *F*-statistics, *p*-values, instantaneous chi-squares, and instantaneous *p*-values for Granger tests between pro-Navalny and pro-government time series. The direction of causality is depicted by arrows.

Before the demonstrations on 21 April, the Russian authorities had reportedly detained approximately 17,600 protesters (OVD-Info, 2021a). In February, a journalist was also imprisoned for retweeting information about the first street rally (Treisman, 2021). In addition, it is worth considering the potential impact of increased pressure from the Russian authorities on Twitter itself. As early as March, they started slowing down Twitter due to the platform's noncompliance with content removal requests (Roskomnadzor, 2021). The throttling was lifted only in May, underscoring the relevance of this pressure to the protest movement and providing another potential explanation behind the noticeable decrease in protest-related tweets. Notably, no large-scale demonstrations have been held since the 21 April protests.

To probe deeper into the relationship between pro-Navalny and pro-government tweeting activity, we performed the Granger causality test on each protest phase. The results presented in Table 1 show no substantial Granger causality between the pro-Navalny and pro-government series, except for two instances. First, a statistically significant Granger-causality effect is observed from pro-government tweets to pro-Navalny tweets during the protests on 14 February ($p < .05$). Second, during the protests on 21 April, pro-Navalny tweets Granger-caused pro-government tweets ($p < .01$). In addition, instantaneous relations between the series emerged during the first and second periods between protests and during the protests on 21 April. According to these results, we can conclude that pro-Navalny tweeting activity provoked the counteraction by pro-government accounts during the protests on 21 April, just as it might have done during the 23 January protests as well ($p = .05$). Furthermore, the pro-government communication on Twitter evoked pro-Navalny accounts' reaction only during the courtyard protests on 14 February.

Given that the latter finding might appear counterintuitive, we inspected the content of popular tweets posted on 14 February. Pro-government accounts were promoting posts supporting veterans and contained related hashtags, particularly highlighting an

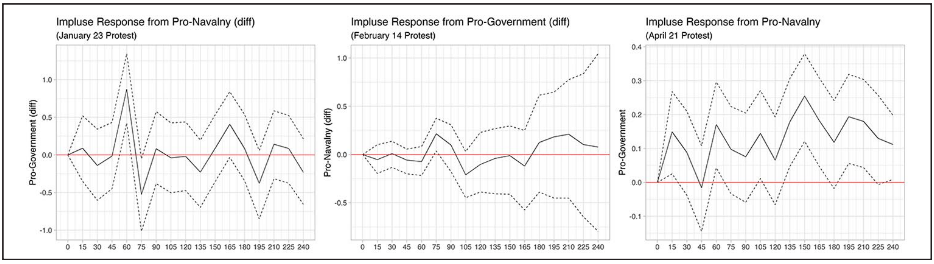

Figure 2. Impulse response function.
The plots illustrate the impact over time of a 1-unit change in the series mentioned in the title on the series mentioned on the y-axis.

offline pro-government flashmob where school children used drones to write the hashtag #ProtectVeterans in the sky. Interestingly, we noticed a decrease in tweets expressing support for the riot police compared to the January protests. This shift could be attributed to the nature of the 14 February protests, organized as short flashmobs in scattered neighborhoods. Navalny supporters extensively shared tweets with the hashtag #LoveIsStrongerThanFear, sharing who they met in courtyards and how many flashlights they saw that evening. Considering that the pro-government flashmob slightly preceded the pro-Navalny one and began to be tweeted earlier that day, the Granger causality might reflect this temporal sequence of the events. Alternatively, protesters may have tried to diminish the visibility of pro-government hashtags by promoting #LoveIsStrongerThanFear.

We further investigated the series that exhibited statistical significance in the Granger causality test using the IRF analysis. We found that an increase in the volume of pro-Navalny tweets resulted in a statistically significant increase in the volume of pro-government tweets during the 23 January and 21 April protests and that the pro-government camp had a statistically significant positive impact on the pro-Navalny tweeting activity during the 14 February protest (see Figure 2). The main statistically significant effect occurred in all these cases after about 1 hour.

Examining the log-transformed and differentiated series allows us to interpret the IRF scale in percentage changes. In the case of the 23 January protests, the pro-Navalny series had to be differentiated twice to achieve stationarity. However, the IRF is almost identical to that resulting from a single-differencing procedure. The most significant impact on the pro-government activity becomes evident after 1 hour, where a 1% increase in the pro-Navalny tweets determines a 0.88% increase in the pro-government tweets. In the case of the 14 February protests, we differentiated the series once. Here, a 1% increase in pro-government tweets resulted in approximately a 0.25% increase in tweets supporting Navalny after about 75 minutes. Finally, we analyzed the log-transformed series for the 21 April protests that did not require differencing. In this case, one log-unit increase in tweets in favor of Navalny corresponded to approximately 0.1–0.2 log units in pro-government tweets in the following period, indicating a 10–20% increase. This effect appears to be more sustained over time: a significant peak is recorded at about 1 hour, after 15 minutes, and after 2.5 hours, indicating a sustained reaction of pro-government actors to tweets supporting Navalny. In all three cases, wide confidence intervals pose challenges in drawing definitive conclusions regarding potential effects immediately before and after the statistically significant peaks.

Pro-Navalny and pro-government coordination patterns

To address RQ2, we measured the synchronization, centralization, and modularity of coordinated networks. Our findings show that accounts supporting Navalny were more likely to engage in coordinated retweeting and did so more quickly than those who supported the government. For instance, 9793 accounts in the pro-Navalny network retweeted a tweet within 1 second, compared to only 243 in the pro-government network. This finding aligns with the overall more intense activity of pro-Navalny accounts, who tweeted more on average ($M=7.80$, $SD=64.08$) than those supporting the regime ($M=3.37$, $SD=7.78$).

Moreover, the pro-Navalny coordinated networks demonstrated higher synchronization and co-retweet frequencies than the pro-government ones across all time intervals (see Figure 3). However, this difference becomes less pronounced for larger time intervals and lower numbers of co-retweets. Interestingly, approximately 50% of users in both camps co-retweeted at least once within a 3-minute time frame. In addition, around 25% of all users co-retweeted twice or more within half an hour. This suggests that the pro-government networks were comparably coordinated within larger periods.

Figure 4 presents the *centralization* scores for the pro-Navalny and pro-government coordinated networks, indicating a low degree of centralization within 1 second of co-retweeting, with a noticeable decline in the centralization scores of these networks as the number of co-retweets increased. However, when the time for co-retweeting increased, coordinated networks from both sides tended to have more centralized structures.

The pro-Navalny networks showed higher centralization as the number of co-retweets increased, particularly at 10 seconds and eight to ten co-retweets, indicating a star-shaped structure. Conversely, the pro-government networks tended to be more centralized as the number of co-retweets declined, though showing an unstable pattern. Compared to pro-Navalny networks, most pro-government networks had lower centralization scores, indicating moderate centralization with the highest values at seven to nine co-retweets within 30 minutes or 1 hour. Therefore, the pro-government networks demonstrated a looser and slower organization of coordination compared to the highly centralized pro-Navalny networks.

Figure 5 shows the modularity values for the coordinated networks of the pro-Navalny and pro-government camps within different time frames and co-retweet numbers. In general, modularity values above 0.3 indicate a significant community structure in a network (Clauset et al., 2004). Correspondingly, we found that pro-government coordinated networks were more effectively organized into distinct communities than those supporting Navalny. The latter consistently showed higher modularity only when the number of co-retweets was equal to one. Interestingly, as the duration and frequency of co-retweeting increased, the number of coordinated pro-government networks with high modularity also increased. This finding suggests that the coordinated networks backing

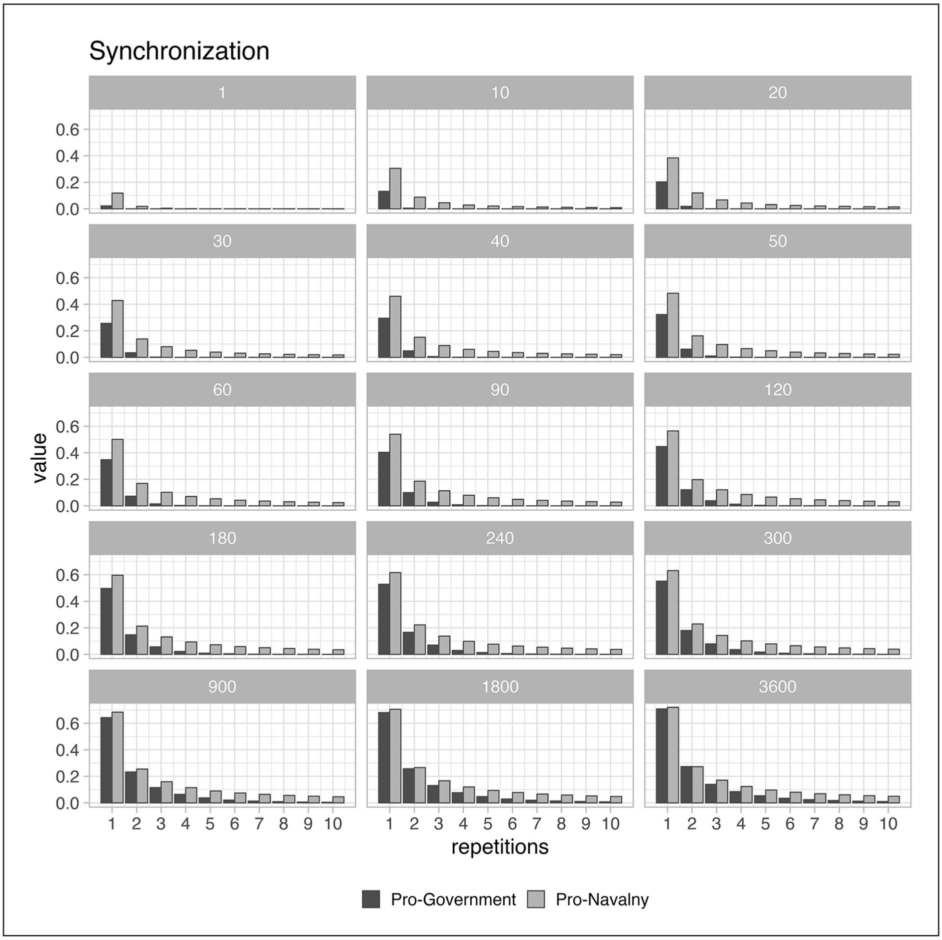

Figure 3. Synchronization of coordinated users.
The *y*-axis displays the proportion of coordinated users, the *x*-axis displays the number of co-retweets, and each plot represents a time window in seconds.

the government contained multiple groups actively engaged in co-retweeting. Notably, such networks tended to co-retweet fewer posts within shorter time windows, for example, only one tweet within 1 second, two within 10 seconds, three within half a minute, and so on. This could indicate a deliberate delay in their co-retweeting activity.

Discussion

This study conceptualized and examined protest-related connective action and counter-action in an authoritarian context, applying Bennett and Segerberg's (2013) connective action framework. Using time series analysis, we first analyzed the dynamic interplay between Twitter activities of pro-Navalny and pro-government accounts and how these

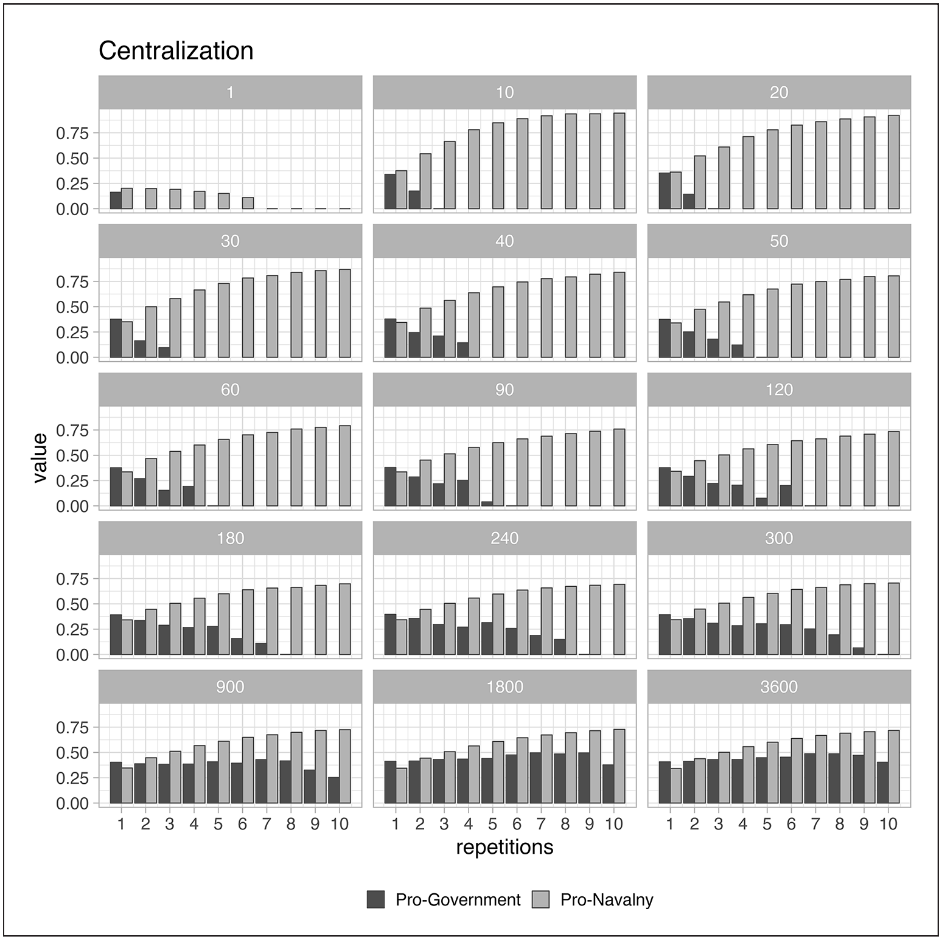

Figure 4. Centralization of coordinated networks.
The *y*-axis displays the centralization score (0 = completely decentralized, 1 = completely centralized), the *x*-axis displays the number of co-retweets, and each plot represents a time window in seconds.

relate to offline protest events. Second, we focused on coordination as a main structural characteristic of connective action (Bennett and Segerberg, 2013) and analyzed co-retweeting networks of pro-Navalny and pro-regime accounts, thereby distinguishing three related but distinct aspects of coordination: synchronization, centralization, and modularity.

Our time series analyses showed that protesters were significantly more active on Twitter during the first month of the movement than their pro-regime counterparts. Nonetheless, pro-regime accounts appeared to intensify their Twitter activity and sustain it right before the protest days. Interestingly, tweeting activity from both sides declined following the February protests and remained low before the April protests

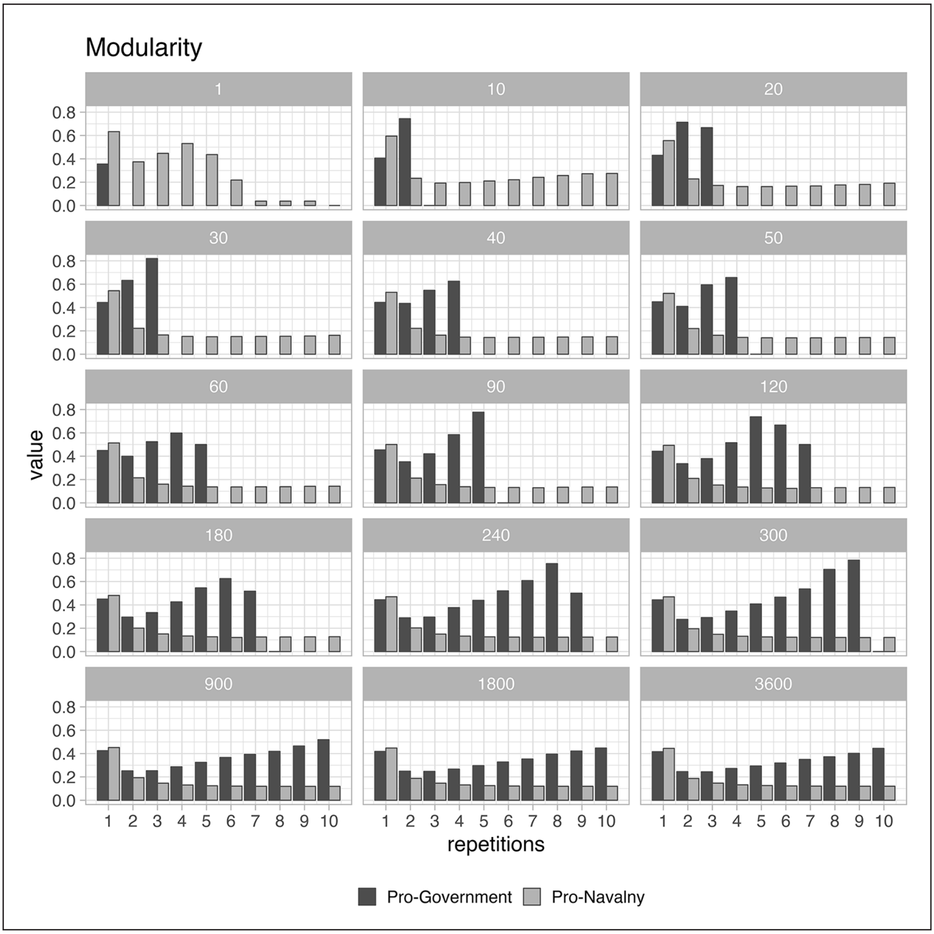

Figure 5. Modularity of coordinated networks.
The *y*-axis displays the modularity values, the *x*-axis displays the number of co-retweets, and each plot represents a time window in seconds.

despite their large nationwide scale. This trend may be attributed to factors such as the throttling of Twitter by the Russian government from March to May as a form of censorship (Xue et al., 2021), alongside increased government repression targeting both participants in the initial protests (OVD-Info, 2021a) and those who retweeted about them (e.g. Treisman, 2021).

Moreover, Granger-causality tests revealed important insights into the relationship between pro-government and pro-Navalny tweeting activities during three key protests. Specifically, pro-Navalny tweets led to a pro-regime reaction during the protests on 21 April and likely on 23 January. Conversely, during the 14 February protests, pro-government tweets preceded counteractions from pro-Navalny accounts. This inconsistent directionality points to the dynamic and reactionary nature of digital protest and repression. One possible explanation for the counterintuitive finding from the 14 February protests is that protesters may have been aware of the pro-regime strategies and responded with their own narrative, thereby resisting the impact of repression. Another plausible explanation involves the inclusive personal action frame, which has the potential to amplify connective action (Bennett and Segerberg, 2013), as exemplified by the motto #LoveIsStrongerThanFear. It transcends specific political affiliations, focusing instead on a personalized and shareable idea. When such ideas evolve into memes, they become widespread and powerful tools of protest communication in Russia (Shomova, 2022).

Another possible explanation is an attempt to diminish the visibility of pro-regime tweets posted just before the courtyard protests of 14 February. Notably, this counteraction occurred a month before the government implemented Twitter throttling. It is conceivable that in the absence of such infrastructure-level censorship, the regime's strategy was to target digital communication on Twitter itself, provoking counteraction. This aligns with the research of Keremoğlu and Weidmann (2020), highlighting the significance of content control by authoritarian regimes, especially when restricting access to specific content is not feasible. This could also explain the pro-regime counteraction to the protests on 23 January and 21 April, where the pro-government accounts responded to an inundation of pro-protest tweets by stirring up the riot police online. These findings suggest an extensive use of organizationally brokered networks by the authorities and their supporters in collective repression against protesters. Adopting Bennett and Segerberg's (2013) lens, we can view pro-regime actors as organizations in the foreground who utilize their high resource brokerage to suppress dissent, with social media tactics serving as a supplementary tool.

Further findings regarding the structural characteristics of the coordinated networks support the above classification. Specifically, the pro-regime coordinated networks show robust modularity, characterized by numerous groups actively involved in co-retweeting. Similar observations were reported by Keller et al. (2020), who attributed this tendency to the division of astroturfing agents into smaller teams operating in various Internet cafes to orchestrate a disinformation campaign. Moreover, most hashtags in pro-government tweets, such as #WorkBrothers, #ProtectVeterans, or #NavalnyIsEnemyOfRussia, focus on a shared goal or purpose. According to Bennett and Segerberg (2013), such collective action framing is commonly used for organizationally brokered collective action and uncommon for connective action.

Interestingly, similar characteristics can be recognized in the tweeting activity of pro-Navalny accounts. They have a high proportion of coordinated accounts that use hashtags with abstract collective action frames, such as #freeNavalny, #RussiaWillBeFree, or #RussiaWithoutPutin. Nevertheless, unlike the pro-government coordinated networks, the pro-Navalny networks were not organized into many distinct communities. Instead, they show little modularity for multiple co-retweets across all time windows, except for the 1-second interval. Co-retweeting more than once in such a short time frame might suggest the presence of some (semi-)automated accounts (Keller et al., 2020) used for scaling up the protests, while the lower modularity observed across the remaining intervals indicates that pro-Navalny tweeters did not coordinate in multiple stable groups.

Nevertheless, their coordinated networks tend to be highly centralized, indicating the presence of influential Twitter accounts whose tweets are often retweeted. Following a post hoc investigation, we identified several opposition figures and ordinary users among the most retweeted accounts, with the official account of Navalny's team (@ teamnavalny) playing a central role. This observation aligns with the findings of Toepfl (2018), who noted that leading Russian opposition activists transitioned from connective to collective action during the 2011–2013 protests. Our study corroborates this trend, suggesting that a decade was insufficient for the opposition to reduce centralization and ignite a more successful crowd-enabled action in 2021. Therefore, even though both protest and repression on Twitter resemble organizationally enabled connective action, they still tend to adhere to the logic of conventional collective action due to highly organized coordination.

The findings of this study should be interpreted with caution. While Twitter API v2 provided comprehensive data on online protest and repression, the sampled data omits deleted tweets. Deletions can occur for various reasons, including account protection, suspension, or self-censorship, which are significant considerations when studying repressive regimes. To assess the magnitude of this issue, we checked the compliance status of our data sets 2 years after the start of initial data collection.[6] The results showed that 81.3% of all tweets in the pro-Navalny data set were still available through the Twitter API, while 77.44% of the pro-government tweets could still be collected. In the pro-Navalny data, unavailability was due to protected user accounts (6.07%), deleted tweets (4.93%), deleted user accounts (3.97%), account suspension (3.54%), and deactivated user accounts (0.15%). In the pro-government data, unavailability was due to account suspension (14.49%), protected user accounts (3.27%), deleted user accounts (3.22%), deleted tweets (1.31%), and deactivated user accounts (0.26%). Therefore, a suggestion for future research in authoritarian contexts is to collect data as close to the dates of protest events as possible while monitoring and documenting observations on trending hashtags.

Another limitation concerns hashtag overlaps within pro-Navalny and pro-government tweets, affecting less than 1% of our data. Due to the insignificant number, we opted to retain these tweets. Similarly, we included potentially irrelevant or contradictory tweets, as their proportion was minimal in the validation samples and did not introduce systematic bias to the results (see Supplemental Appendix B). Finally, as Twitter does not reflect the general population, the findings cannot be generalized beyond the platform (Lazer et al., 2021). Future research could collect data from multiple social media platforms to draw conclusions about a broader population. In the Russian context, exploring potential avenues could involve examining platforms, such as Telegram and YouTube, both widely used and accessible (as of January 2024).

Notwithstanding these limitations, our study contributes to the understanding of digital protest and repression dynamics within authoritarian regimes. It posits that digital repression may manifest as a genuine countermovement while being organized through coordinated accounts. Assuming it is organized by the government, the regime initially allows digital protest, likely for the reasons articulated in Toepfl (2020), such as gathering information about society, opposition, and grievances. Simultaneously, it endorses counteractions to online protests to increase the perceived risks of participating in street demonstrations. This is achieved by covert information channeling (Earl et al., 2022), expressed in demonstrative support for the government. While such actions may sometimes succeed in marginalizing oppositional voices (Spaiser et al., 2017), our findings indicate that in 2021, the opposition was notably more vocal on Twitter compared to their pro-regime counterparts. When the opposition effectively counters repression, the government may employ heavier measures, such as throttling and arrests. To gain a deeper understanding of this interplay, we recommend conducting further studies to explore the dynamics of connective action and counteraction, focusing on the role of coordinated networks. The R application developed in this study serves as a convenient tool for identifying and studying such networks.[7] Moreover, with reliable access to data, a similar approach can be employed for platform auditing or investigating a range of social media phenomena beyond authoritarian regimes. This includes but is not limited to examining social movements and countermovements, polarization dynamics, and the spread of disinformation.

Acknowledgements

The authors would like to acknowledge Petro Tolochko for his valuable contribution to data validation and the anonymous reviewers for their excellent feedback.

Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

ORCID iDs

Aytalina Kulichkina https://orcid.org/0000-0001-5948-9881
Nicola Righetti https://orcid.org/0000-0002-9257-5113
Annie Waldherr https://orcid.org/0000-0001-7488-9138

Supplemental material

Supplemental material for this article is available online. The scripts for data collection and analyses are accessible at: https://osf.io/8769b/.

Notes

1. Data were collected using the academictwitteR R package (Barrie and Ho, 2021).
2. We used the tidyverse (Wickham et al., 2019) and the gridExtra (Auguié and Antonov, 2017) R packages.
3. We used the tseries (Trapletti and Hornik, 2022) R package.
4. We used the vars (Pfaff, 2008a, 2008b) R package.
5. We used the igraph (Csardi and Nepusz, 2006) R package.
6. We used twCompliance (Schatto-Eckrodt, 2022) R package.
7. The R application was later expanded into the R package CooRTweet, accessible at https://cran.r-project.org/web/packages/CooRTweet/index.html.

References

Auguié B and Antonov A (2017) GridExtra: miscellaneous functions for “grid” graphics (software). Available at: https://CRAN.R-project.org/package=gridExtra

Baklanov A (2021) Incoming, trolls Thousands of suspicious accounts target Russian opposition figures linked to Alexey Navalny on Instagram. *Meduza*, 20 January. Available at: https://meduza.io/en/feature/2021/01/20/incoming-trolls (accessed 19 May 2021).

Barberá P, Wang N, Bonneau R, et al. (2015) The critical periphery in the growth of social protests. *PLoS ONE* 10(11): e0143611.

Barrie C and Ho JC (2021) AcademictwitteR: an R package to access the Twitter academic research product track v2 API endpoint. *Journal of Open Source Software* 6(62): 3272.

Bastos MT, Mercea D and Charpentier A (2015) Tents, tweets, and events: the interplay between ongoing protests and social media. *Journal of Communication* 65(2): 320–350.

Bennett WL and Segerberg A (2013) *The Logic of Connective Action: Digital Media and the Personalization of Contentious Politics*. Cambridge: Cambridge University Press.

Blondel VD, Guillaume JL, Lambiotte R, et al. (2008) Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment* 2018: 10008.

Bruns A, Highfield T and Burgess J (2013) The Arab Spring and social media audiences: English and Arabic Twitter users and their networks. *American Behavioral Scientist* 57(7): 871–898.

Clauset A, Newman ME and Moore C (2004) Finding community structure in very large networks. *Physical Review E* 70(6): 066111.

Csardi G and Nepusz T (2006) The Igraph software package for complex network research. *Inter Journal, Complex Systems* 1695(5): 1–9.

Earl J and Kimport K (2011) *Digitally Enabled Social Change: Activism in the Internet Age*. Cambridge, MA: MIT Press.

Earl J, Maher TV and Pan J (2022) The digital repression of social movements, protest, and activism: a synthetic review. *Science Advances* 8(10): eabl8198.

Feldstein S (2021) *The Rise of Digital Repression: How Technology Is Reshaping Power, Politics, and Resistance*. Oxford: Oxford University Press, Incorporated.

Freelon D, McIlwain C and Clark M (2018) Quantifying the power and consequences of social media protest. *New Media & Society* 20(3): 990–1011.

Giglietto F, Righetti N and Rossi L (2020a) CooRnet. *Detect coordinated link sharing behavior on social media (software)*. Available at: https://ora.uniurb.it/handle/11576/2675493#.YYJjJi8w0lI

Giglietto F, Righetti N, Rossi L, et al. (2020b) It takes a village to manipulate the media: coordinated link sharing behavior during 2018 and 2019 Italian elections. *Information, Communication & Society* 23(6): 867–891.

Glazunova S (2022) *Digital Activism in Russia: The Communication Tactics of Political Outsiders*. Cham: Springer.

Graham T and QUT Digital Observatory (2020) Coordination network toolkit (software). Available at: https://doi.org/10.25912/RDF_1632782596538

Granger CWJ (1969) Investigating causal relations by econometric models and cross-spectral methods. *Econometrica* 37(3): 424–438.

Hafner CM and Herwartz H (2009) Testing for linear vector autoregressive dynamics under multivariate generalized autoregressive heteroskedasticity. *Statistica Neerlandica* 63(3): 294–323.

Hellmeier S (2016) The dictator's digital toolkit: explaining variation in internet filtering in authoritarian regimes. *Politics and Policy* 44(6): 1158–1191.

Howard PN, Woolley S and Calo R (2018) Algorithms, bots, and political communication in the US 2016 election: the challenge of automated political communication for election law and administration. *Journal of Information Technology & Politics* 15(2): 81–93.

Jost JT, Banaji MR and Nosek BA (2004) A decade of system justification theory: accumulated evidence of conscious and unconscious bolstering of the status quo. *Political Psychology* 25(6): 881–919.

Keller FB, Schoch D, Stier S, et al. (2020) Political astroturfing on Twitter: how to coordinate a disinformation campaign. *Political Communication* 37(2): 256–280.

Kendall-Taylor A, Frantz E and Wright J (2020) The digital dictators: how technology strengthens autocracy. *Foreign Affairs* 99: 103–115.

Keremoğlu E and Weidmann NB (2020) How dictators control the internet: a review essay. *Comparative Political Studies* 53(10–11): 1690–1703.

Killick R and Eckley IA (2014) Changepoint: an R package for changepoint analysis. *Journal of Statistical Software* 58(3): 1–19.

King G, Pan J and Roberts ME (2017) How the Chinese government fabricates social media posts for strategic distraction, not engaged argument. *American Political Science Review* 111(3): 484–501.

Lazer D, Hargittai E, Freelon D, et al. (2021) Meaningful measures of human society in the twenty-first century. *Nature* 595(7866): 189–196.

Lonkila M, Shpakovskaya L and Torchinsky P (2021) Digital activism in Russia: the evolution and forms of online participation in an authoritarian state. In: Gritsenko D, Wijermars M and Kopotev M (eds) *The Palgrave Handbook of Digital Russia Studies*. Cham: Springer; Palgrave Macmillan, pp. 135–153.

Lotan G, Graeff E, Ananny M, et al. (2011) The revolutions were tweeted: information flows during the 2011 Tunisian and Egyptian revolutions. *International Journal of Communication* 5: 1375–1405.

Lütkepohl H (2006) *New Introduction to Multiple Time Series Analysis*. New York: Springer.

Marwick A and Lewis R (2017) Media manipulation and disinformation online. *Data & Society Research Institute*. Available at: http://www.chinhnghia.com/DataAndSociety_MediaManipulationAndDisinformationOnline.pdf

Morozov E (2011) *The Net Delusion: The Dark Side of Internet Freedom*. New York: PublicAffairs.

Navalny A (2021) Tweets [@ navalny] [Twitter profile]. *Twitter*. Available at: https://web.archive.org/web/20210117205444/https://twitter.com/navalny (accessed 16 April 2024).

News and Trending Topics in Russia (2021) News and trending topics in Russia. *50 Trends*. Available at: http://www.50trends.com/ru/ (accessed 17 January–31 December 2021).

OVD-Info (2021a) «Freedom of assembly guaranteed, 17,600 protesters detained»: rebuttal of the Russian response to UN inquiry on widespread assembly rights violations. Available at: https://en.ovdinfo.org/comments-russian-governments-submission-response-joint-enquiry-special-procedures?_gl=1*1j4dx5v*_ga*MjQ0NDU4NjYyLjE2ODU0NzE2OTc.*_ga_J7DH9NKJ0R*MTY4NTQ3MTc0MC4xLjEuMTY4NTQ3MjY2MS40MC4wLjA.#2 (accessed 5 March 2023).

OVD-Info (2021b) The crackdown on peaceful protests in January — February 2021 in Russia. Available at: https://en.ovdinfo.org/winter-2021-supression-en?_gl=1*1v1wr5f*_ga*MTA0MTM3NDMxNC4xNjc4MDM1MTM5*_ga_J7DH9NKJ0R*MTY3ODAzNTE0MC4xLjEuMTY3ODAzNjQ1Mi4zNy4wLjA.#1 (accessed 5 March 2023).

Pfaff B (2008a) *Analysis of Integrated and Cointegrated Time Series with R*. 2nd ed. New York: Springer.

Pfaff B (2008b) VAR, SVAR and SVEC models: implementation within R package vars. *Journal of Statistical Software* 27(4): 1–32. Available at: https://www.jstatsoft.org/v27/i04/

Poupin P (2021) Social media and state repression: the case of VKontakte and the anti-garbage protest in Shies, in Far Northern Russia. *First Monday* 26(5): fm.v26i5.11711. Available at: https://doi.org/10.5210/fm.v26i5.11711 (accessed 7 January 2022).

Roskomnadzor (2021) Roskomnadzor took measures to protect Russian citizens from the influence of illegal content. Available at: https://web.archive.org/web/20210310070927/https://rkn.gov.ru/news/rsoc/news73464.htm (accessed 5 March 2023).

Schatto-Eckrodt T (2022) Hidden biases—the effects of unavailable content on Twitter on sampling quality. In: Jünger J, Gochermann U, Peter C, et al. (eds) *Grenzen, Probleme und Lösungen bei der Stichprobenziehung*. Köln: Herbert von Halem Verlag, pp. 178–195.

Shin Y (2017) *Time Series Analysis in the Social Sciences: The Fundamentals*. Oakland, CA: University of California Press.

Shomova S (2022) “Hieroglyphs of protest”: Internet memes and the protest movement in Russia. *Problems of Post-Communism* 69(3): 232–241.

Sobol L (2021) Интересно, кто-то обратился в твиттер (возможно, кремлевские боты массово развлекаются), что я якобы хочу покончить жизнь самоубийством. *Twitter*. Available at: https://archive.vn/8AzXD#selection-3017.0-3017.250 (accessed 15 March 2023).

Spaiser V, Chadefaux T, Donnay K, et al. (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.

Statista (2023) Number of Twitter users in Russia from 2017 to 2021(in millions). Available at: https://www.statista.com/forecasts/1144549/twitter-users-in-russia (accessed 21 April 2024).

Steinert-Threlkeld ZC, Mocanu D, Vespignani A, et al. (2015) Online social networks and offline protest. *EPJ Data Science* 4(1): 1–9.

Team Navalny (2021) Tweets [@teamnavalny] [Twitter profile]. *Twitter*. Available at: https://web.archive.org/web/20210117184244/https://twitter.com/teamnavalny (accessed 16 April 2024).

Toepfl F (2018) From connective to collective action: Internet elections as a digital tool to centralize and formalize protest in Russia. *Information, Communication & Society* 21(4): 531–547.

Toepfl F (2020) Comparing authoritarian publics: the benefits and risks of three types of publics for autocrats. *Communication Theory* 30(2): 105–125.

Tornes A (2021) Enabling the future of academic research with the Twitter API. Available at: https://developer.twitter.com/en/blog/product-news/2021/enabling-the-future-of-academic-research-with-the-twitter-api (accessed 7 March 2023).

Trapletti A and Hornik K (2022) Tseries: time series analysis and computational finance (software). Available at: https://cran.r-project.org/web/packages/tseries/index.html

Treisman R (2021) Russian journalist jailed after court deems joking retweet a “call to protest..” *NPR*, 4 February. Available at: https://www.npr.org/2021/02/04/964129810/russian-journalist-jailed-after-court-deems-joking-retweet-a-call-to-protest (accessed 6 January 2022).

Tucker JA, Nagler J, Metzger MM, et al. (2016) Big data, social media, and protest: foundations for a research agenda. In: Alvarez RM (ed.) *Computational Social Science: Discovery and Prediction*. Cambridge: Cambridge University Press, pp. 199–224.

Tufekci Z and Wilson C (2012) Social media and the decision to participate in political protest: observations from Tahrir Square. *Journal of Communication* 62(2): 363–379.

Wasserman S and Faust K (1994) *Social Network Analysis: Methods and Applications*. Cambridge: Cambridge University Press.

Weidmann NB and Rød EG (2019) *The Internet and Political Protest in Autocracies*. New York: Oxford University Press.

Wickham H, Averick M, Bryan J, et al. (2019) Welcome to the tidyverse. *Journal of Open Source Software* 4(43): 1686.

Xu X (2021) To repress or to co-opt? Authoritarian control in the age of digital surveillance. *American Journal of Political Science* 65(2): 309–325.

Xue D, Ramesh RS, Valdik SS, et al. (2021) Throttling Twitter: an emerging censorship technique in Russia. In: *Proceedings of the 21st ACM internet measurement conference* (eds D Levin and A Mislove), New York, 2–4 November, pp. 435–443. New York: ACM.

Author biographies

Aytalina Kulichkina is a Research Associate and PhD candidate at the University of Vienna. Her research focuses on the intersection of political communication, social media, and computational methods in social science.

Nicola Righetti is an Assistant Professor at the University of Urbino Carlo Bo. His research focuses on computational and statistical methods and the study of digital political communication.

Annie Waldherr is Professor of Computational Communication Science at the University of Vienna. She studies the changing structures and dynamics in today's digitized public spheres, combining computational and conventional empirical methods.