

---

# A Synchronized Action Framework for Detection of Coordination on Social Media

Thomas Magelinski, Lynnette Hui Xian Ng and Kathleen M. Carley

---

**Abstract.** The study of the coordinated manipulation of conversations on social media has become more prevalent as social media’s role in amplifying misinformation, hate, and polarization has come under greater scrutiny. We discuss how successful generalized coordination detection algorithms could be used to reinforce existing power imbalances, such as those between marginalized groups and government agencies. We propose an alternative method of identifying manipulation—*detecting synchronized actions*—which reduces this risk. We further consider how responsible coordination detection may be carried out by analyzing synchronized actions. We propose a synchronized action framework for detecting automated coordination by constructing and analyzing multi-view networks. We validate our framework by examining a large Twitter dataset surrounding the Reopen America conversation from 2020. We first discover three simple coordinated campaigns, and then investigate synchronized actions between users discussing the protests that could be consistent with covert coordination. This task is far more complex than examples evaluated in prior work, which demonstrates the need for our multi-view approach. Next, we identify a cluster of suspicious users and detail the activity of three members. These three users amplify protest messages using the same hashtags at very similar times, though they all focus on different states. This analysis highlights the potential usefulness of coordination detection algorithms in investigating amplification, as well as the need to carefully and responsibly deploy such tools.

---

## 1 Introduction

Coordinated influence operations pose unique threats to social cybersecurity. While people can be trained to recognize accounts that are likely bots, coordinated users may be inconspicuous until compared to other accounts. Covert coordination can play a role in inciting protests. Indeed, recent political protests in the United States have been orchestrated by external actors unbeknownst to their attendees (Weiss 2018). Inorganic coordination between communities threatens the social fabric of society and has the potential to trigger offline violence. Examples include coordinated efforts to

propagate ISIS's extremist ideas and coordinated efforts to share false stories about racially-motivated attacks during the Black Panther movie event in February 2018 (Benigni, Joseph, and Carley 2017; Babcock, Beskow, and Carley 2018).

Network-based approaches are particularly useful for detecting sets of coordinating users. These approaches draw connections between users when they both take an action that could be considered coordinated. Drawing these connections between users within an entire dataset results in a user-to-user "coordination network." The more often that pairs of users exhibit behavior that could be coordinated, the stronger the connection between them in the coordination network. While it is very difficult for analysts to sift through hundreds of thousands of tweets to find patterns between users, they can easily identify these patterns when the coordination network is plotted; dense clusters of users point to sets of suspicious users, who have been exhibiting potentially coordinated behavior among themselves. When the social media dataset is too big, the full coordination network may be difficult to visualize. Network-based clustering algorithms can help point analysts in the right direction by automatically detecting sets of users that are densely connected.

Social media users may coordinate along a number of behaviors. For example, a group of users may want to push a collection of hashtags or a collection of URLs. Prior work on using network-based methods to detect coordinated users studied one method of coordination at a time: one process to study hashtags and another to assess URLs; the results then had to be combined. Second, Further, current methods of identifying synchronized behavior between users miss as many as half such instances.

We developed a multi-view network-based synchronized action framework that resolves both problems. Multi-view networks allow analysts to study several types of coordinated behavior at the same time; they no longer have to run multiple network analyses and combine the results themselves. The synchronized action aspect of our approach fixes the error found in previous models that resulted in the loss of up to 50% of the synchronized connections between users for the action type being studied, such as hashtag synchronization. We consider actions that can be directly inferred from social media data (e.g., URL, hashtag and mentions), and exclude those that require an initial processing step to uncover, such as similar text or images (Pacheco et al. 2020; Ng, Cruickshank, and Carley 2021).

While the automated detection of coordinated accounts remains an important problem, some solutions have the potential to reinforce existing power imbalances between the state and civil organizers. Since prior research has largely overlooked the effect on this imbalance, we discuss it below. We argue that responsible coordination detection can be performed by analyzing highly synchronized actions, a tactic that is misleading to the general population and unlikely to be used by genuine political organizers due to the very short timescale. While political organizers rely on synchronized efforts, the type of synchronization we discuss here is on a scale that is unlikely to be seen by genuine political organizers. This is further discussed in Section 3.

We apply our approach, which we call the multi-view synchronized action framework, to a Twitter dataset surrounding the Reopen America protests of 2020 to uncover coordination through hashtags, URLs, and mentions. This approach is validated by the discovery of three "templated" coordinated campaigns, in which tweets within the campaign follow a simple template (e.g., they might all use the same two hashtags while tweeting a specific URL). However, covert coordination is likely to be more complex and nuanced than templated examples and those seen in prior work, as users will take steps to hide their tracks. By focusing on the part of the dataset that more directly discusses the Reopen America protests, we identify a cluster of users engaging in more sophisticated

synchronization. Within this cluster, we highlight three accounts that exhibit synchronized behavior in support of the protests that is consistent with a coordinated effort to promote them.

## 2 Background

The field of social cybersecurity is exploding in interest as social media's role in fanning the flames of civil unrest has reached the forefront of public discussion (Carley et al. 2018). This is especially true following the Washington, D.C. protests and subsequent breach of the nation's Capitol on January 6, 2021, which was fueled by voter fraud mis/disinformation both on and off social media (Berlinski et al. 2021). While "fake news" on social media has led to many real-world consequences, many unanswered questions remain (Grinberg et al. 2019), including: How can we discover sets of users who are coordinating to manipulate online discussions?

One of the closest lines of inquiry is bot detection, which seeks to identify automated users on social media (Ferrara et al. 2016; Cresci 2020). Common bot-detection mechanisms are either (1) *feature based*, using account details such as tweet frequency and content quality (Beskow and Carley 2018; Davis et al. 2016; K.-C. Yang et al. 2020), or (2) *graph based*, which rely on an account's communication links (Cao et al. 2012; Skornjakov, Turdakov, and Zhabotinsky 2018; Magelinski, Beskow, and Carley 2020). While feature-based approaches scale better to large datasets, they only consider accounts in isolation.

The bot-detection algorithms that are closest to coordinated-account detection are DeBot and CopyCatch. DeBot (Chavoshi, Hamooni, and Mueen 2016) finds correlated accounts on Twitter based on the temporal activity of user tweets, but does not account for tweet content, an essential component of detecting coordinated behavior. CopyCatch (Beutel et al. 2013) was built to detect coordinated attempts to inflate the number of "likes" on Facebook pages. While it can be generalized to identify synchronized actions, it is designed to detect coordination within a single burst of time, rather than the long-term coordination we can detect using a network-based approach.

The gap between bot detection and coordinated activity has recently been recognized. In an attempt to close this gap, Pacheco et al. (2020)'s framework for analyzing coordination involves (1) defining a coordinated action type (e.g., hashtag usage), (2) creating a user-behavior bipartite network that encodes the connections drawn between users and the hashtags they have used, and (3) constructing a coordination network by folding the user-behavior network into a user-user network; network links (i.e., the number of times a pair of users used the same hashtag) indicate the strength of the shared behavior. Typically, their networks are sparse enough that the resulting networks consist of small sets of users who are well connected to each other, but not to anyone else. These disconnected sets of users are called network components, and the study of these suspicious groups of users is known as "component analysis" (Wasserman, Faust, et al. 1994).

When there are too many users in a component to effectively analyze, Nizzoli et al. (2020) suggests using a more principled procedure that involves network community detection algorithms, which separates users into densely connected sets when applied to the coordination network. Analysts can then start by investigating the densest clusters.

Vargas, Emami, and Traynor (2020) advanced the use of network statistics to predict coordinated activity on Twitter, defining coordination behaviors in terms of retweets, co-tweets, hashtags and so forth. They find that users coordinate by synchronizing along a number of behaviors, for example when they retweet specific tweets or use

the same hashtags or text. While groups can coordinate along multiple dimensions, our study is the first to investigate these dimensions of coordination at the same time. We do so using multi-view network analysis (Kivelä et al. 2014). Multi-view networks, sometimes referred to as multi-layer networks, are an active area of research due to their flexibility to model complex interactions (Aleta and Moreno 2019; Boccaletti et al. 2014; De Domenico et al. 2013; Zhang et al. 2017; Hutchinson et al. 2019). While a regular network simply draws edges between pairs of users, a multi-view network draws *different types* of edges between users. This information is stored separately, so analysts can observe that a pair of users is connected in more than one way. For example, an analyst can evaluate a pair of users in a multi-view network to see how often they use the same hashtag *and* how often they use the same URLs.

The multi-view approach has often been used to develop clustering algorithms, which are relevant to the problem of coordination detection (Bickel and Scheffer 2004; Yang and Wang 2018; Cruickshank 2020). Multi-view clustering succeeds because it can create clusters of items that are not all connected in the same way. In the context of coordination detection, a group of users may seek to increase the visibility of a hashtag and of a URL. The group may take a non-uniform approach in which some users tweet the hashtag, some users tweet the URL, and others tweet both. Using the classic approach, these users will be placed into different clusters, and analysts may or may not realize that they have some overlap and should be combined. With a multi-view network, these users can all be placed in the same cluster.

Simply converting existing synchronized-action techniques into a multi-view network is insufficient. The most commonly used method of detecting synchronized actions fails to identify many instances of such actions (Pacheco et al. (2020), Pacheco, Flammini, and Menczer (2020), and Weber and Neumann (2020)). This approach first breaks the timeline into discrete time windows, and is thus referred to as the discrete-window approach. Any instances of two users performing the same action within the same time window is then recorded as a synchronized action. Yet this approach can generate a problem. For example, if we set the time window size to 5 minutes, breaking up the timeline could result in one window from 1:00pm to 1:05pm, and the next from 1:05pm to 1:10pm. If user 1 tweets #reopen at 1:04pm, and user 2 tweets #reopen at 1:06pm, the actions fall into separate time windows despite being only 2 minutes apart, a smaller differential than our set time of 5 minutes. If actions are uniformly spread out over time, this windowing approach will miss 50% of the synchronized actions. Overlapping time windows could be used to obtain these edges, yet some instances will be double counted, leading to a different type of inaccuracy.

We seek to resolve these methodological issues in two ways: (1) by defining coordination detection as a *multi-view* network problem, permitting the detection of coordination across different types of actions and (2) by providing an algorithm to record *all* instances of synchronized action in each view.

### 3 The Ethical Implications of Coordination Detection

Given the increasing interest in social media manipulation, it is important to be able to detect coordinated actors such platforms. However, to the best of our knowledge, prior studies of coordinated behavior have not explored the ethical implications of this problem further than acknowledging that it is difficult to distinguish between “good” and “bad” coordination. We push this line of inquiry further and employ Kalluri’s framing to investigate how successful coordination detection may shift power (Kalluri 2020). To limit shortsightedness, we consider unrealized “successful” coordination detection (Russell

2019) in the context of protests—a common area of study for social media coordination (Magelinski and Carley 2020; Ng, Cruickshank, and Carley 2021; Steinert-Threlkeld et al. 2015; Enikolopov, Makarin, and Petrova 2020).

Ideally, coordination detectors can identify meaningful sets of users who are cooperating in some way on social media. Social media is an effective tool for alternative protest reporting, since it gives the public the power to directly report events as they experience them (Enikolopov, Makarin, and Petrova 2020; Hermida and Hernández-Santaolalla 2018), allowing them to find and organize support for a cause, which sometimes escalates into offline protests (Steinert-Threlkeld et al. 2015).

Law enforcement and government agencies around the world are increasingly engaging in the surveillance of social media (Mateescu et al. 2015; Patton et al. 2017; Williams et al. 2013; Qin, Strömberg, and Wu 2017), including policing online actions as well as protests and mobilization (Dencik, Hintz, and Carey 2018). Social media information about protests and political dissent has also been used to proactively disrupt organized protests by preemptively arresting key actors in the organizer networks (Dencik, Hintz, and Carey 2018; Swain 2013). Such use of social media against the public has led many to conclude that the power of social media has shifted away from the public towards law enforcement and government agencies (Mateescu et al. 2015). Previous demonstrations and riots have been used to justify further developing the technologies driving this power shift (Williams et al. 2013).

*Unrestricted coordination detection* thus runs the risk of exacerbating imbalances in power against the public, especially for minority communities that already face a large power imbalance (Owen 2017; Qin, Strömberg, and Wu 2017; Patton et al. 2017). Yet *specialized* detectors that shift power in another way may not do so.

A *synchronized action* occurs when two users take the same action at the same time. We assess the *degree of synchronization* based on the timescale in which actors operate. While normal coordination may mobilize people to talk about a certain topic on a certain day, users performing highly synchronized coordination take the same actions within minutes of one another. Such actions can artificially inflate the popularity of narratives, ideas, or topics, which can be used to manipulate people. We focus on detecting highly synchronized actions in an effort to shift power away from groups attempting such manipulation, towards ordinary users to increase their online information.

## 4 Synchronized Action Framework

We define an “action” as any measurable behavior a user exhibits that can be localized to a particular point in time, such as a user tweeting “#reopen.” “Action types” are more general sets of potential actions, such as hashtag usage; tweeting #reopen, #liberate, or #openup are different actions of the hashtag type.

Action types are used to define the views of the multi-view coordination network. For each action type, users are connected to each other based on the strength of their synchronization among actions within their type. If hashtag usage is an action type of interest, there will be a hashtag view in the coordination network, connecting users based on how strongly their hashtag usage was synchronized.

The definition of synchronized action can be relaxed by considering either similar *actions* or similar *times*. DeBot, for example, searches for users who have correlated tweet times (Chavoshi, Hamooni, and Mueen 2016). Placing their approach within our framework, the action *and* action type is simply to tweet, giving a one-view coordination network. The measure of synchronization used is temporal correlation. Thus, the choice of action types,

synchronization measure, and network analysis approach fully defines the method.

We use a relaxed *same-action, similar-time* definition of synchronized action. We first consider three standard action types—hashtag, URL usages, and mention usages—which results in a three-view network. We consider highly synchronized actions occurring within a 5-minute window of each other using the method described in Section 4.2. We observe similar results when the time window is set from 1–10 minutes. The 5-minute window size was selected on the basis that it is short enough that users will be unlikely to fall within the same time window multiple times by chance, but long enough that purposefully coordinating actors could accomplish synchronization through means other than direct automation. Shorter time windows also make analysis easier, as there are fewer edges for the analyst to explore. For our data, 5 minutes seemed to strike the right balance, though this window size may need to be adjusted depending on the specifics of other datasets it is applied to.

We then use multi-view clustering to identify the cluster with the highest density, and investigate central nodes within this dense cluster (Cruickshank 2020) by visually inspecting the coordination network to understand which actors are displaying synchronized behaviors, and how frequently. Strong connections, visualized as thick lines between nodes, are studied by reading the specific tweets that contained the synchronized actions.

Lastly, we consider three “higher-order” actions, or combinations of actions: hashtag and URL, URL and mention, and hashtag and mention. These actions are only counted if both actions occur in the same tweet. While such actions result in subsets of the original actions, the direct connection between the content (hashtag or URL) and target (mention) makes discovered coordinated campaigns easy to analyze under social cyber-security frameworks, like the BEND framework, which gives analysts guidance about the implications of targeted cyber-operations (Carley 2020).

#### 4.1 Multi-View Coordination Network

A multi-view coordination network is an  $L$ -layer network  $G = \{V, E, \mathcal{L}\}$  where  $\mathcal{L}$  is the set of view indices  $\{1, 2, \dots, L\}$ .  $V = \{V^1 \cup V^2 \dots \cup V^L\}$ , where  $V^i$  denotes the set of nodes in layer  $i$  of the network.  $E = \{E^1 \cup E^2 \dots \cup E^L\}$ , where  $E^i$  denotes the set of edges in layer  $i$  of the network. Specifically, nodes  $V$  are Twitter users, and edges  $E$  are connections between users based on coordination actions. Each view  $L^i \in \mathcal{L}$  represents an action type  $i$ . Standard action types are single instances like a common hashtag between two tweets. Higher-order action types are a combination of standard action types: edges for the (hashtag-URL) action type are formed when two tweets contain a common hashtag and URL.

In each view  $L^i$ , the action type  $i$  can have multiple actions,  $A^i = \{A_1^i, A_2^i, \dots, A_k^i\}$ . For example, the hashtag action type encompasses many actions: #reopenNY, #reopenTX or #reopenPA. An edge  $e_{x,y} \in E^i$  represents the presence of coordination action type  $i$  between users  $x$  and  $y$  within a time window  $t$ . The network can contain directed or undirected edges, or both. Weighted edges  $e_{x,y}$  have a weight value  $w_{x,y} > 0$ , which represents the strength of coordination between the two users. Here, we consider undirected weighted edges.

In our first example, a multi-view coordination network is formed with three layers, or three standard action types:  $L = \{\text{hashtag, URL, mention}\}$ . For example, if users  $x$  and  $y$  post tweets with the hashtag “#reopenNY” within  $t$ , an edge  $e_{x,y} \in E^1$  is drawn. In our second example, we form a multi-view coordination network with three higher-order action types:  $L = \{(\text{hashtag-URL}), (\text{URL-mention}), (\text{hashtag-mention})\}$ . For this work,

we set  $t = 5$  minutes.

## 4.2 View Construction

The connections within each view represent the frequency of synchronized actions between two users within the corresponding action type. For example, the hashtag view encodes how users are synchronizing their hashtag usage. As discussed above, the “fixed window” approach employed in previous studies fails to identify many connections within these views. In this section, we outline a scalable approach that uncovers all of the connections.

A naive approach to obtain all links is to sort tweets chronologically, move a *sliding* window across them, and then count the instances of shared behavior. This approach does not scale. It presents an  $O(N^2)$  time complexity, where  $N$  is the number of tweets that require pairwise comparison, meaning that a dataset of 1,000 tweets will take 100 times longer than one that contains 100 tweets,  $(\frac{1000}{100})^2 = 100$ . This is unsustainable for analyzing datasets with millions of tweets.

We only need to compare tweets with the same action type. If we are considering hashtags, we only need to determine when two tweets with the same hashtag are posted within the same window. Thus, we first group the tweets by their action, and run the sliding window only within these groups. The procedure is still  $O(N^2)$ , however the  $N$  has been greatly reduced, improving the runtime and making it feasible to analyze datasets with millions of tweets. The number of calculations needed is directly proportional to the number of edges in the coordination network, making this approach the optimal solution. This implies that more specific action types are more scalable. For example, if the action type is the most general action possible, simply tweeting,  $N$  is still the set of all tweets. A very specific action type like tweeting both a hashtag and a URL will greatly reduce  $N$ .

We also observe that the shorter the time window, the more quickly the approach will run. Shorter windows will entail fewer coordinated actions and will therefore run faster. Thus, the “highly coordinated” behavior that we set out to focus on in Section 3 is built into our approach.

When a sliding window is used, as opposed to the time segments employed in previous works, the weighting of connections becomes non-trivial. This is particularly important for coordination, because erroneously strong connections can make normal users look suspicious. In a static window, the strength of the connection between two users should correspond to the minimum number of instances between them in the window. For example, if user A is in the window twice, while user B is in the window 100 times, the connection should only be of strength 2. Otherwise, spammers will be erroneously strongly connected to normal users. In a sliding window, this is more complex because the total number of interactions between the two users is not directly known. To account for this characteristic, we use a heuristic: the user with a lower presence in the current time window is likely to have fewer connections overall. Based on this heuristic, we only draw connections from the user in question to others if the user has a lower presence in the time window than the other potential connecting user. This procedure gives similar edge weights to the simple fixed-window example.

Our approach is theoretically optimal, but empirical data on how long the methods take to run is still useful. Applying our approach to a dataset with more than 3.6 million accounts took an average of 13.02 seconds for each of the three studied views, resulting in a total time of 39.06 seconds.

### 4.3 Multi-View Network Analysis

After constructing the multi-view network, we cluster it using the multi-view modularity clustering technique (Cruickshank 2020), which seeks to combine different views of the data to produce a better clustering output. We focus on the cluster with the highest density since this implies many coordinated actions.

We then identify the nodes with the highest centrality in terms of total degree, eigenvector centrality, and community–hub modularity vitality (Wasserman, Faust, et al. 1994; Magelinski, Bartulovic, and Carley 2021). Total degree is of particular importance, as it identifies the nodes with the most coordinated connections between users who are the most likely to truly be coordinating.

### 4.4 Action Type Selection

We begin by selecting standard action types on Twitter: sharing hashtags, sharing URLs, and mentioning other users. However, these action types can lead to false-positive connections, particularly in the case of hashtags, where some users were observed to be strongly connected but tweeting opposing views. As such, singular action types of hashtag and URL usage may be too general.

We therefore combine the standard action types into higher-order actions of order 2. For example, for the (hashtag-URL) action type, we consider tweets that contain both a hashtag *and* a URL. Tweets that contain multiple hashtags and URLs are separated into all possible (hashtag, URL) tuples. The notion of higher-order action leads to more specific behavior, with which makes erroneous user connections less likely. This also generates a scalability advantage. The definition of very specific behavior within a short sliding time window leads to few but meaningful edges in the coordination network graph, making the network analysis easier, as prior work on coordination detection has shown.

## 5 Results

In this section, we present the results of coordination detection on a Twitter dataset collected around the Reopen America Protests of 2020. Ideally, our approach would be applied to an existing public dataset with ground-truth labels for users who are coordinating. However, ground-truth datasets for coordination may all have a different definition of “ground truth,” as the definition of coordinating accounts is not completely agreed upon; definitions range from similar individual behaviors such as linguistic properties Addawood et al. 2019, to anomalously high levels of coincidental behavior (Weber and Falzon 2021), to participation in the promotion of a common narrative (Pacheco et al. 2020). Further, our method can only detect a very specific type of coordination—highly synchronized coordination—which may not be present in benchmark datasets.

Instead, we demonstrate the validity of our approach by illustrating its effectiveness in detecting “templated” campaigns. Tweets in such campaigns follow a very specific pattern and are thus easy to verify as part of a coordinated effort. While the three templated campaigns identified are not engaging in the conversation around the protests, they demonstrate that our approach can detect coordinated users.

The coordination network of users who *are* engaging in the protest discussion is far more complex. Within this network, we uncover a set of users who may have coordinated to promote pro-reopen messages during the protests.

## 5.1 Data

The Reopen America protests occurred across the United States from April to September 2020 and had the goal of lifting the COVID-19 safety restrictions affecting daily work and recreation. Some protesters disregarded social distancing and mask protocols, and confronted reporters during the protests (Tracy 2020; BBC 2020; Shepherd and Balingit 2020). The similarity of protest organizing materials, dates, and phrasing across states raised questions about coordination, making this an interesting dataset to examine (Stanley-Becker and Romm 2020; Chandler 2020).

We used Twitter Search API to collect a large corpus of tweets from April 1, 2020 to June 22, 2020. We searched for keywords or hashtags containing “openup,” “reopen,” “operationgridlock,” or “liberate.” All US state abbreviations were appended to each of the search terms (e.g., “liberateNY”) to collect the Reopen America protest for each state. These terms were selected based on early observations of hashtags surrounding the issue. They capture tweets discussing the protest as well as other discussions—a common issue with keyword-based data collection. Nevertheless, we find coordinated activity in many discussions both related to the protests and not. The dataset contains roughly 3.6 million unique users and 9.9 million tweets.

## 5.2 Validation through the Discovery of Templated Campaigns

We first demonstrate the effectiveness of our framework by describing how it identified three templated campaigns. We apply our approach to simultaneously consider coordination with hashtags, mentions, and URLs. Using this analysis, we discover the 7News network. Our higher-order action analysis then identifies campaigns that rely on a single higher-order action (e.g., mentioning the same users while posting the same URL in the same tweet) rather than a multitude of standard actions. Based on this observation, these higher-order action networks are studied using one view at a time. We find notable campaigns in the URL-mention and hashtag-mention networks, which contained the Mexico Without Plastics Campaign and the Reopen ICT campaign, respectively.

### 5.2.1 7News Network

Before conducting the multi-view clustering, we perform a visual network analysis. We first observe many small components in the multi-view network, each of which contains publicly affiliated accounts such as news networks. The strongest link in the network (averaged across views) connects two accounts that are formally affiliated, providing initial face validity.

The best-connected component is the Australian 7News network, shown in Figure 1. These accounts consistently rank highest in both degree centrality and community–hub modularity vitality. The network is highly coordinated along all action types:  $A^{\text{hashtag}}$ ,  $A^{\text{URL}}$  and  $A^{\text{mention}}$ .

A visual network analysis of Figure 1 clearly reveals that this is a coordinated group, as every account has taken synchronized actions with every other account, and they have used all the potential methods of synchronization. This highly coordinated network results in a complete graph formation structure among the 7News account, which is the simplest templated coordination campaign. The “template” used by these accounts is that they all tweet the same (or nearly identical) tweet at the same time, which amplifies the message to reach a broader audience.

The tweets often mention the reporter, using the story’s URL, and adding hashtags like #7News. They appear in our dataset for their coverage of “reopen” events, mostly

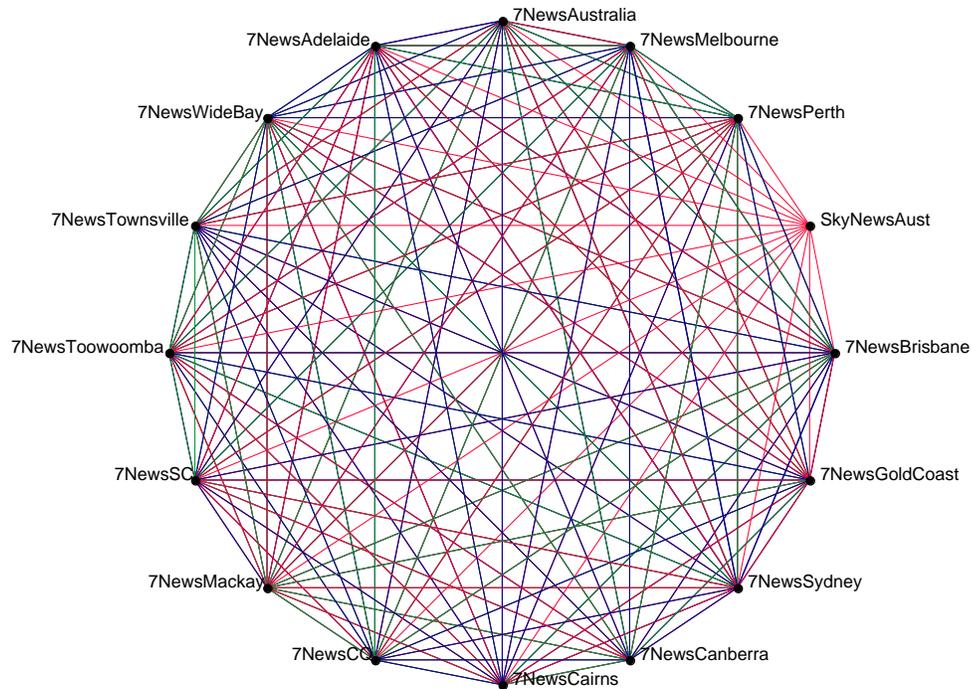


Figure 1: 7News component of the multi-view coordination network. Link width corresponds to coordination strength. Blue, red, and green lines correspond to URL, mention, and hashtag coordination, respectively.

focusing on the eventual reopening of Australia’s borders. The coordination between local accounts broadens their geographic coverage, and reinforces the popularity of their URLs and hashtags.

It is evident that these accounts coordinate along all action types. Eleven accounts of the 7News network shared the news on the expected date for reopening Queensland’s borders. This message was shared several times within the time span the message was active, resulting in a weight of 5–7 between any two accounts in all views:

July 10 is firming as the most likely date for Queensland’s borders to reopen, but it could happen even sooner if the number of COVID-19 cases across Australia remains low. <https://7news.com.au/> @Bianca\_Stone #7NEWS

Shared by: 7NewsGoldCoast, 7NewsSydney, 7NewsCairns, 7NewsCQ, 7News-GoldCoast, 7NewsMackay, 7NewsSC, 7NewsToowoomba, 7NewsTownsville, 7NewsiwdeBay, 7NewsMackay

Time span message was active: 2020-06-14 20:58:16 to 2020-06-14 21:02:42

In another example of coordination within this network, 12 accounts shared the same message a total of 180 times, with the same time stamp. We infer that the news network may have used automated methods to push out the same message through different account channels; the target audience of each is various cities in Australia. However, we did not find the same text being tweeted by all 15 discovered 7News accounts, which we attribute to Twitter’s API sampling process. Twitter’s API returns a randomized sample of the tweets that meet the collection criteria, so it is likely that the dataset is missing some of the 7News tweets.

New South Wales ski resorts will be allowed to reopen slopes on June 22 after being given the green light by the state government. <https://t.co/3JOp14F6v6> @MyleeHogan #7NEWS <https://t.co/36bdIm48T7>

Shared by: 7NewsCQ, 7NewsGoldCoast, 7NewsCairns, 7NewsMackay, 7NewsSC, 7NewsMelbourne, 7NewsToowoomba, 7NewsTownsville, 7NewsWideBay, 7NewsSydney, 7NewsAustralia, 7NewsBrisbane

Initial time shared: 2020-05-18 01:57:36

Lastly, we observe that the account @SkyNewsAust is connected to this network by the hashtag action only, suggesting that this account leveraged the hashtag coined by the 7News network. We posit that this helped its tweets gain visibility by latching onto hashtags propagated by several high-visibility news accounts.

### 5.2.2 Mexico Without Plastics Campaign

Investigating the strongest connections in the  $L^{\text{URL-mention}}$  view of the higher-order network uncovers the “Mexico Without Plastics” campaign to ban single-use plastics. This campaign was sponsored at least in part by Greenpeace Mexico, which hosted a petition calling for Mexico’s Senate to reform the country’s waste-management laws to ban single-use plastic. The petition garnered over 275,000 signatures. Figure 2 displays the ego network of the most active account. It illustrates that a few coordinating actors were key leaders in promoting the campaign. These coordinating actors are “hubs” placed in the center of the network, which have many coordinated links to those on the outskirts. In this campaign, the URL is the petition, while the mention is either the official account of the Mexican Senate, or its members. The campaign is present in our dataset due to the keyword “libérate,” as in “free yourself from plastic.”

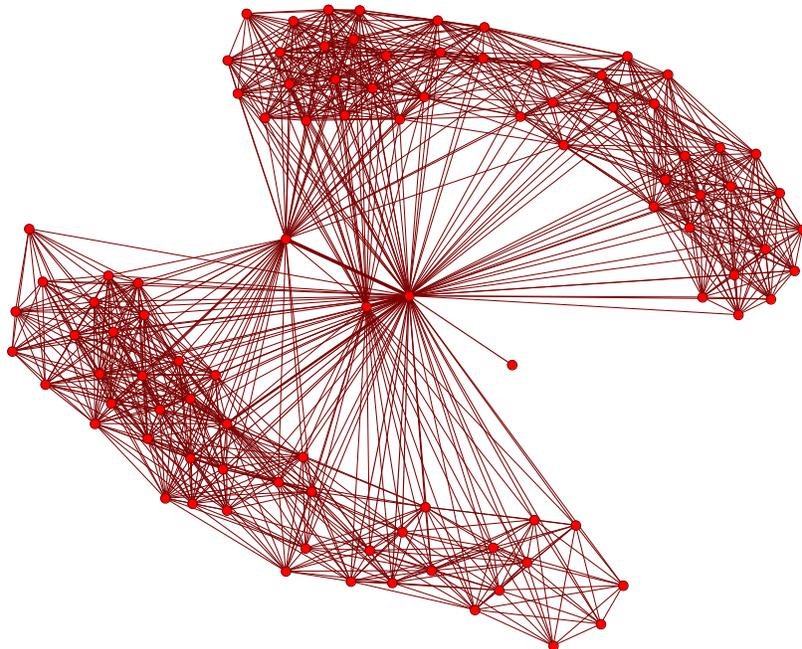


Figure 2: Ego of the URL-mention coordination network of the most-connected users in the Mexico Without Plastics Campaign.

Table 1 presents two examples of tweets from two authors of the URL-mention coordina-

tion network.

Table 1: Examples of synchronized tweets from two authors the URL-Mention coordination network in the Mexico Without Plastics campaign. Author names are redacted to maintain author privacy.

Text 1	Text 2
Las prohibiciones de plásticos de un solo uso no ponen en riesgo nuestra salud, la contaminación plástica sí. Recordemos al @senadomexicano y a @RaulBCCue que no hay salud en un planeta enfermo. #MéxicoSinPlásticos <a href="https://t.co/oiFZ8HfRdw">https://t.co/oiFZ8HfRdw</a> 2020-06-18 14:53:06	¡Podemos vivir sin plásticos, pero no sin un planeta sano! Exijo a @senadomexicano y a @RaulBCCue no dejen que la industria plástica se aproveche de la pandemia para impulsar el uso de desechables fuera de uso médico. Quiero un #MéxicoSinPlásticos <a href="https://t.co/xwC6xd7Ji6">https://t.co/xwC6xd7Ji6</a> 2020-06-18 15:00:15
Las prohibiciones de plásticos de un solo uso no ponen en riesgo nuestra salud, la contaminación plástica sí. Recordemos al @senadomexicano y a @RaulBCCue que no hay salud en un planeta enfermo. #MéxicoSinPlásticos <a href="https://t.co/oiFZ8HfRdw">https://t.co/oiFZ8HfRdw</a> 2020-06-18 14:53:06	Exijamos al @senadomexicano y a @RaulBCCue que se apoye una legislación a favor del medio ambiente, no de la industria plástica. ¡Que no se eche abajo el avance de la ciudadanía con la prohibición de plásticos desechables! #MéxicoSinPlásticos <a href="https://t.co/xwC6xd7Ji6">https://t.co/xwC6xd7Ji6</a> 2020-06-18 15:00:49

### 5.2.3 Campaign to Open ICT for Indian Tax Inspectors

Lastly, the  $L_{\text{hashtag-mention}}$  view of the higher-order network is studied. We observe a dense cluster of users lacking the “hubs” observed in the Mexico Without Plastics campaign, suggesting a lack of centralized leadership. These users are campaigning in India to grant inspectors in the tax office the ability to have internal transfers. The ICT, or Inter-Commissionerate Transfer, allows Indian employees to be transferred from one state to another. This has become increasingly important in the coronavirus pandemic, allowing employees to transfer to other states and reunite with their families during the lockdown period. Thus, the users are coordinating with hashtags like #DepressionKillsOpenICT and #CGSTinspectorsWantICT, while mentioning Prime Minister Modi, government officials, and news outlets. The campaign appears in our dataset due to phrase “reopen ICT”.

Table 2 displays instances of synchronized tweets between two sets of authors. In these tweets, the users are desperately calling for the authorities to reopen ICT as it constitutes a primary source of income and provides welfare for their families. These tweets also tag politicians (@PMOIndia) and organizations (@cbic\_india), likely in hopes of gaining their attention. The community leverages each other’s tweets to amplify them.

### 5.3 Coordinated Amplification of Reopen Protests

While the previous examples demonstrate the ability of our approach to detect coordinated campaigns, the campaigns themselves are templated examples. They are intentionally public facing, clearly coordinated, and result in strong network structures such as disconnected components and extremely strong edges. The coordinated activity surrounding the Reopen America protests is less clear cut and more interesting, illustrating the difficulty of detecting covert coordination.

Table 2: Examples of synchronized tweets from two authors the hashtag-mention coordination network in the Campaign to Open ICT. Author names are redacted to maintain author privacy.

Text 1	Text 2
Author pair 1	
#OpenICTinCBICnow Please don't deny us the natural right of being with our family. Kindly reopen ICT in CBIC for Inspector Cadre. @narendramodi @cbic_india @nsitharama @FinMinIndia @DoPTGoI @IR-SAssociation @PMOIndia @ianuragthakur @thewire_in @ThePrintIndia @TheLallantop	#OpenICTinCBICnow ICT is a genuine demand of inspector. Many of us are leaving hard earned job bcoz of ICT ban. Kindly reopen ICT. @narendramodi @cbic_india @nsitharama @FinMinIndia @DoPTGoI @IR-SAssociation @PMOIndia @ianuragthakur @thewire_in @ThePrintIndia @TheLallantop
2020-05-24 07:31:00	2020-05-24 07:35:54
Author pair 2	
Sir/Madam Please reopen ICT for Inspectors in CBIC.Our family is suffering very much due to this. #RestoreICTinCBIC @cbic_india @Anurag_Office @ianuragthakur @PMOIndia @nsitharaman @sgaiceia @ravishndtv @RahulGandhi @DGHRDPOLICY @aahtak	#RestoreICTinCBIC We should not let down the mothers among us. ICT is a easy to implement transfer policy. It won't do any harm to anyone. We hope @cbic_india will do their best to reopen ICT for all ASAP. @Anurag_Office @ianuragthakur @PMOIndia @nsitharaman @sgaiceia
2020-06-01 14:31:11	2020-06-01 14:34:48
Sir/Madam Please reopen ICT for Inspectors in CBIC.Our family is suffering very much due to this. #RestoreICTinCBIC @cbic_india @Anurag_Office @ianuragthakur @PMOIndia @nsitharaman @sgaiceia @ravishndtv @RahulGandhi @DGHRDPOLICY @aahtak	#RestoreICTinCBIC CBIC Board Members are not looking after welfare of their staff and working as per their whims and fancies.They are not following DOPT guidelines. Plz reopen the ICT for Inspectors. @cbic_india @Anurag_Office @ianuragthakur @PMOIndia @nsitharaman @sgaiceia
2020-06-01 14:31:11	2020-06-01 14:35:07

We find strong connections within the densest cluster of the first-order multi-view coordination network, as found through multi-view modularity clustering. We study some of the strongest connections in this cluster. The most central node in terms of total degree centrality is @FedUpUSA, a self-proclaimed media account committed to bringing “truth about what is really happening, as opposed to the fodder that is shown in the mainstream media.” Its website has been inactive for 2 years, while its Twitter account remains active.

The multi-view ego network of @FedUpUSA in Figure 3 demonstrates its complexity. The coordinated activity around the Reopen protests is far murkier than that the campaigns discussed above. The lack of simple coordination is demonstrated by the absence of any obvious strongly connected clique, which indicates either a lack of coordination or a more sophisticated strategy. Thus, network measures and account content analysis become even more important than they were in the simple examples seen in the validation section and in prior studies.

Since @FedUpUSA is strongly connected to a number of accounts, we focus on the two that are most strongly synchronized with it, though anecdotally many of the associated accounts display similar behavior. Both users appear to be private citizens or have

protected accounts, so we refer to them as Accounts 2 and 3.

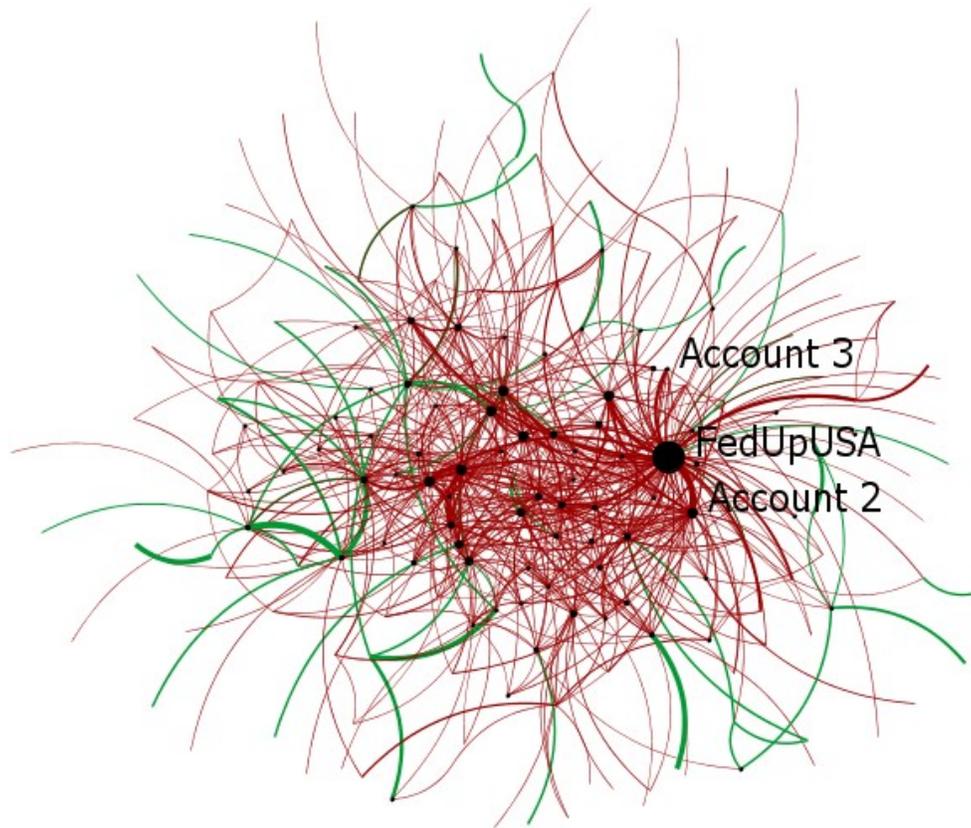


Figure 3: FedUpUSA multi-view coordination ego network. Red lines denote mentions; green lines correspond to hashtag coordination. Link width corresponds to the number of synchronized actions between the two nodes. Nodes are sized by their total degree.

All accounts were highly active during the protests, but each had a different regional focus: @FedUpUSA focused on Michigan, Account 2 on Minnesota, and Account 3 on North Carolina. All accounts pushed hashtags related to the protests; the most popular were #OperationGridlock and #EndTheShutdown, each of which was used over 40 times combined between the users, not counting retweets. @FedUpUSA was particularly focused on hashtag usage, often tweeting pictures of the Michigan protests with up to 10 hashtags while mentioning news outlets and public officials. Account 2 tweeted and retweeted supportive statements about the Michigan protests, before focusing on Minnesota. The Minnesota tweets included the times and locations of one protest, livestream information, and supportive statements from public figures such as the president. Account 3 did not focus on the protests themselves, but on sending the “reopen” message directly to the North Carolina governor and other officials through mentions. They frequently used hashtags such as #ReopenNC and #VoteCooperOut. They also focused on reopen “successes,” such as a court order allowing religious organizations to meet.

All three accounts produced unique, coherent content and engaged with other users through intelligible replies to their posts. These activities are extremely difficult to automate, and are inconsistent with the expected automated behaviors that bot-detection methods leverage (Davis et al. 2016; Beskow and Carley 2018; K.-C. Yang et al. 2020; Cresci 2020). These differences highlight the distinction between coordination detection and bot detection. While these accounts had strong overlap in hashtag and mention

usage, they did not share any of the same URLs, highlighting the difference in coordination strategies between groups such as the 7News organization, which relied heavily on URLs.

Table 3 lists examples of the synchronized tweets, which illustrate the interaction between @FedUpUSA and other accounts in campaigning for the reopening of America. Trump is a major target of these campaigns: users are trying to convince him to take further actions. Unlike the templated campaigns, these accounts engage in a combination of tweets with nuanced text and logical arguments and tweets simply promoting hashtags and URLs. This mixed approach highlights the difficulty of identifying these messages as synchronized without the use of automated tools.

## 6 Limitations

This study suffers from a number of limitations. First, our notion of coordination is highly limited. More importantly, without knowledge of private affiliations or plans, we cannot truly say that particular accounts are coordinating. Instead, our methods point to suspicious users that are taking the same actions at the same time. It is entirely possible that the users we found to be synchronized during the reopen protests did not in fact work together behind the scenes, so we make no claims about direct account affiliation when this is not made clear by the users themselves. However, we note that regardless of affiliation, their impact is similar since they pushed similar messages at similar times.

Figure 2 illustrates that synchronization-based networks create network chains, in which many users are weakly connected because their tweets happen to occur at a similar time to many others. This highlights the study's next limitation—the challenge associated with differentiating coordinated efforts from organically emerging support. We tested a weighting scheme that down-weighted popular actions, which successfully minimized “false positive” connections, though it obscured the overall network structure. Thus, future work should develop weighting schemes that minimize inadvertent connections and retain the coordination structure.

Further, while the coordinated actors we discovered appear to be spreading the same messaging, this is not guaranteed, especially under single-view analysis. Hashtags in particular can serve multiple roles, so a hashtag-only analysis could be connecting users with opposing views on a hashtag that are talking at the same time (L. Yang et al. 2012). Further, we may be inadvertently detecting “hashtag hijacking”—the practice of using popular hashtags to increase viewership of one's tweets. However, we believe this is mitigated through the other action type connections and the focus on strong links. Future work leveraging tweet text and possibly sentiment would be even stronger.

The collection strategy also creates a bias in the tweets for  $L^{\text{hashtag}}$  and against  $L^{\text{URL}}$ . By collecting tweets using keywords, we are restricted to finding URLs that either contain the terms themselves or are posted along with one of the keywords. The significantly higher number of connections in the hashtag view compared to the link view may affect the quality of the final clusters, since the current techniques in multi-view clustering work best under similar view statistics.

Lastly, our dataset is a primarily English dataset from Twitter due to the chosen keywords. While we observe coordination among English speakers (Reopen protests) and Spanish speakers (plastics campaign), our results are still highly skewed towards users tweeting in English. Future work could explore collecting the same hashtag in multiple languages to build a coordination network across different languages. Additionally, keyword-based

Table 3: Examples of synchronized tweets surrounding @FedUpUSA around the Reopen Protests. Authors names other than @FedUpUSA are redacted to maintain privacy.

Text 1 (written by @FedupUSA)	Text 2
Action: Mentions	
<p>@realDonaldTrump Can someone look into Governor Murphy and his illegal activities? We have an average of 38 COVID 19 patients per hospital in this state, yet no plan to reopen. He is not a leader, all he leads with is fear. Destroying/Destroyed NJ. Look at this beach signage. Where is AG Barr <a href="https://t.co/4ggBeIJh2Y">https://t.co/4ggBeIJh2Y</a></p> <p>2020-05-10 16:37:31</p> <p>@realDonaldTrump @GovMurphy hey you also living on a \$1200 stimulus check or your still getting paid????? #ReOpenNJ</p> <p>2020-05-10 16:41:12</p>	<p>@ElbertForsythe @realDonaldTrump •Models Project Sharp Rise in Deaths as States Reopen • An internal Trump administration report expects about 200,000 daily cases by June. The White House bars coronavirus task force officials from testifying to Congress without approval. <a href="https://t.co/9OuOSkF1LF">https://t.co/9OuOSkF1LF</a></p> <p>2020-05-10 16:37:35</p> <p>@realDonaldTrump I hope so. Our family is out of \$\$, never got #UI even though qualified and never received #stimulus from #IRS. More than 2 months without \$. @GovWhitmer has a tyrannical stranglehold on us. #ReopenMichigan #SendHelp</p> <p>2020-05-10 16:37:35</p>
Action: Hashtags	
<p>#OperationGridlock #OpenMichigan #EndTheShutdown #FreedomForAll #WorkingIsEssential #EducationIsEssential #FreedomOfSpeech #FreedomOfReligion #DriveOutTyranny #RecallWhitmer #MichiganUnitedForLiberty @limbaugh @WayneDupreeShow @FOX2News @wjrradio <a href="https://t.co/Zmuir1OQcl">https://t.co/Zmuir1OQcl</a></p> <p>2020-04-15 16:17:28</p> <p>This is how you deal with tyranny! #OperationGridlock #OpenMichigan #EndTheShutdown #FreedomForAll #WorkingIsEssential #EducationIsEssential #FreedomOfSpeech #FreedomOfReligion #DriveOutTyranny #RecallWhitmer #MichiganUnitedForLiberty <a href="https://t.co/nbnJ2HAipj">https://t.co/nbnJ2HAipj</a></p> <p>2020-04-15 18:20:11</p>	<p>#OperationGridlock #OpenMichigan #EndTheShutdown #FreedomForAll #WorkingIsEssential #EducationIsEssential #FreedomOfSpeech #FreedomOfReligion #DriveOutTyranny #RecallWhitmer All going down right now in front of #WickedWhitmer's fortress. <a href="https://t.co/O25NUxVcop">https://t.co/O25NUxVcop</a></p> <p>2020-04-15 16:18:21</p> <p>Gridlock for 10 miles in all directions. #Lansing #Michigan! #OperationGridlock #OpenMichigan #OpenMI #EndTheShutdown #FreedomForAll #WorkingIsEssential #FreedomOfSpeech #DriveOutTyranny #RecallWhitmer #MichiganUnitedForLiberty <a href="https://t.co/XRfKMntvWy">https://t.co/XRfKMntvWy</a></p> <p>2020-04-15 18:25:44</p>
Action: Hashtag-Mentions	
<p>#OperationGridlock #OpenMichigan #EndTheShutdown #FreedomForAll #WorkingIsEssential #EducationIsEssential #FreedomOfSpeech #FreedomOfReligion #DriveOutTyranny #RecallWhitmer #MichiganUnitedForLiberty @limbaugh @WayneDupreeShow @Local4News @FOX2News <a href="https://t.co/jSJ37v7JnHn">https://t.co/jSJ37v7JnHn</a></p> <p>2020-04-15 15:45:07</p>	<p>@FoxNews @FOX2News @NBCNews @ABC #OperationGridlock #recallwhitmer #RecallGovWhitmer #michiganlockdown #michiganshutdown #lansing #michigan-dersagainstexcessivequarantine it's happening. The streets are locked down. Traffic is backed up for miles. At least 25,000 people are here. <a href="https://t.co/Nyl5KEcZBF">https://t.co/Nyl5KEcZBF</a></p> <p>2020-04-15 15:49:55</p>

collection often retrieves tweets that are not related to the event, which results in our detection of unrelated campaigns, like the Indian Tax campaign.

## 7 Conclusion

Coordination detection can reinforce existing power imbalances. However, we believe coordinated activities on social media can be responsibly detected by analyzing synchronized actions. We therefore propose a multi-view synchronized action framework that facilitates the analysis of coordination behavior across several actions simultaneously. This general framework encompasses behaviors studied in previous works and allows for expansion into others. It is designed to uncover synchronized action within narrow time windows, allowing the detection of highly orchestrated covert campaigns as well as viral, public-facing campaigns, while not investigating coordination on the longer time scales characteristic of grass-roots organizations.

We demonstrate the effectiveness of our approach using the Twitter dataset related to the Reopen America conversation. Our method has optimal scalability and is able to process three types of coordination within the nearly 10 million node dataset in under 40 seconds. We identify three templated coordinated activities: the 7News Network's effort to further its content through locally based accounts tweeting together, the "Mexico Without Plastics" campaign to ban single-use plastics through waste legislation, and the "Reopen ICT" campaign to reopen transfers for Indian tax inspectors. These examples are publicly coordinated and simple to analyze, mirroring those seen in prior work, in which network analysis is straightforward.

These templated coordinated activities contrast starkly with the complexity of synchronized actions observed in discussion of the Reopen America protests. The observation of more nuanced synchronization strategies (i.e. strategies beyond all accounts posting the same content at nearly the same time) could be consistent with covert coordination. With this added complexity, the multi-view approach we introduced becomes necessary. We detect a cluster of suspicious accounts and detail the actions of users who pushed similar hashtags and mentioned the same users at similar times while emphasizing protests in other states individually.

The complexity of the relationships between users that exhibited synchronized actions in the Reopen America protest discussion highlights the need to include even more actions in the multi-action approach, such as tweet text, images, or account creation. Further, analyzing different timescales may be fruitful. The difference in the observed complexity between templated campaigns and more sophisticated synchronization strategies highlights the difficult challenge these methods face in detecting sophisticated actors while avoiding the potential for abuse. Future work is needed on the challenging task of intent discovery, where tools like ours can be integrated into existing workflows. We hope further work exploring these analyses can meet this challenge, and that progress can be made towards sieving out chance synchronicity.

## References

- Addawood, Aseel, Adam Badawy, Kristina Lerman, and Emilio Ferrara. 2019. "Linguistic cues to deception: Identifying political trolls on social media." In *Proceedings of the international AAAI conference on web and social media*, 13:15–25.
- Aleta, Alberto, and Yamir Moreno. 2019. "Multilayer networks in a nutshell." *Annual Review of Condensed Matter Physics* 10:45–62. <https://doi.org/10.1146/annurev-conmatphys-031218-013259>.
- Babcock, Matthew, David M Beskow, and Kathleen M Carley. 2018. "Beaten up on Twitter? Exploring fake news and satirical responses during the Black Panther movie event." In *International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction and Behavior Representation in Modeling and Simulation*, 97–103. Springer. [https://doi.org/10.1007/978-3-319-93372-6\\_12](https://doi.org/10.1007/978-3-319-93372-6_12).
- BBC. 2020. *Coronavirus: Armed protesters enter Michigan statehouse*, May. <https://www.bbc.com/news/world-us-canada-52496514>.
- Benigni, Matthew C, Kenneth Joseph, and Kathleen M Carley. 2017. "Online extremism and the communities that sustain it: Detecting the ISIS supporting community on Twitter." *PloS one* 12 (12). <https://doi.org/10.1371/journal.pone.0181405>.
- Berlinski, Nicolas, Margaret Doyle, Andrew M Guess, Gabrielle Levy, Benjamin Lyons, Jacob M Montgomery, Brendan Nyhan, and Jason Reifler. 2021. "The effects of unsubstantiated claims of voter fraud on confidence in elections." *Journal of Experimental Political Science*, 1–16. <https://doi.org/10.1017/XPS.2021.18>.
- Beskow, David M, and Kathleen M Carley. 2018. "Bot-hunter: A tiered approach to detecting & characterizing automated activity on twitter." In *SBP-BRIMS: International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction and Behavior Representation in Modeling and Simulation*, vol. 8. [http://sbp-brims.org/2018/proceedings/papers/latebreaking\\_papers/LB\\_5.pdf](http://sbp-brims.org/2018/proceedings/papers/latebreaking_papers/LB_5.pdf).
- Beutel, Alex, Wanhong Xu, Venkatesan Guruswami, Christopher Palow, and Christos Faloutsos. 2013. "CopyCatch: stopping group attacks by spotting lockstep behavior in social networks." In *Proceedings of the 22nd international conference on World Wide Web*, 119–30. <https://doi.org/P10.1145/2488388.2488400>.
- Bickel, Steffen, and Tobias Scheffer. 2004. "Multi-view clustering." In *ICDM*, 4:19–26. 2004. Citeseer. <https://doi.org/10.1109/ICDM.2004.10095>.
- Boccaletti, Stefano, Ginestra Bianconi, Regino Criado, Charo I Del Genio, Jesús Gómez-Gardenes, Miguel Romance, Irene Sendina-Nadal, Zhen Wang, and Massimiliano Zanin. 2014. "The structure and dynamics of multilayer networks." *Physics reports* 544 (1): 1–122. <https://doi.org/10.1016/j.physrep.2014.07.001>.
- Cao, Qiang, Michael Sirivianos, Xiaowei Yang, and Tiago Pregueiro. 2012. "Aiding the detection of fake accounts in large scale social online services." In *Presented as part of the 9th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 12)*, 197–210. <https://www.usenix.org/conference/nsdi12/technical-sessions/presentation/cao>.
- Carley, Kathleen M. 2020. "Social cybersecurity: an emerging science." *Computational and Mathematical Organization Theory* 26 (4): 365–81. <https://doi.org/10.1007/s10588-020-09322-9>.

- Carley, Kathleen M, Guido Cervone, Nitin Agarwal, and Huan Liu. 2018. "Social cyber-security." In *International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction and Behavior Representation in Modeling and Simulation*, 389–94. Springer. <https://doi.org/10.1007/s10588-020-09322-9>.
- Chandler, Simon. 2020. *Security researchers say the reopen America campaign is being astroturfed*, April. <https://www.forbes.com/sites/simonchandler/2020/04/24/security-researchers-say-the-reopen-america-campaign-is-being-astroturfed/?sh=49541cc16506>.
- Chavoshi, Nikan, Hossein Hamooni, and Abdullah Mueen. 2016. "Debot: Twitter bot detection via warped correlation." In *IEEE International Conference on Data Mining (ICDM)*, 817–22. <https://doi.org/10.1109/ICDM.2016.0096>.
- Cresci, Stefano. 2020. "A decade of social bot detection." *Communications of the ACM* 63 (10): 72–83. <https://doi.org/10.1145/3409116>.
- Cruickshank, Iain J. 2020. "Multi-view clustering of social-based data." PhD diss., Carnegie Mellon University.
- Davis, Clayton Allen, Onur Varol, Emilio Ferrara, Alessandro Flammini, and Filippo Menczer. 2016. "Botornot: A system to evaluate social bots." In *Proceedings of the 25th international conference companion on world wide web*, 273–74. <https://doi.org/10.1145/2872518.2889302>.
- De Domenico, Manlio, Albert Solé-Ribalta, Emanuele Cozzo, Mikko Kivelä, Yamir Moreno, Mason A Porter, Sergio Gómez, and Alex Arenas. 2013. "Mathematical formulation of multilayer networks." *Physical Review X* 3 (4): 041022. <https://doi.org/10.1103/PhysRevX.3.041022>.
- Dencik, Lina, Arne Hintz, and Zoe Carey. 2018. "Prediction, pre-emption and limits to dissent: Social media and big data uses for policing protests in the United Kingdom." *New Media & Society* 20 (4): 1433–50. <https://doi.org/10.1177/1461444817697722>.
- Enikolopov, Ruben, Alexey Makarin, and Maria Petrova. 2020. "Social media and protest participation: Evidence from Russia." *Econometrica* 88 (4): 1479–514. <https://doi.org/10.3982/ECTA14281>.
- Ferrara, Emilio, Onur Varol, Clayton Davis, Filippo Menczer, and Alessandro Flammini. 2016. "The rise of social bots." *Communications of the ACM* 59 (7): 96–104. <https://doi.org/10.1145/2818717>.
- Grinberg, Nir, Kenneth Joseph, Lisa Friedland, Briony Swire-Thompson, and David Lazer. 2019. "Fake news on Twitter during the 2016 US presidential election." *Science* 363 (6425): 374–78. <https://doi.org/10.1038/s41467-018-07761-2>.
- Hermida, Alberto, and Víctor Hernández-Santaolalla. 2018. "Twitter and video activism as tools for counter-surveillance: the case of social protests in Spain." *Information, Communication & Society* 21 (3): 416–33. <https://doi.org/10.1080/1369118X.2017.1284880>.
- Hutchinson, Matthew C, Bernat Bramon Mora, Shai Pilosof, Allison K Barner, Sonia Kéfi, Elisa Thébault, Pedro Jordano, and Daniel B Stouffer. 2019. "Seeing the forest for the trees: Putting multilayer networks to work for community ecology." *Functional Ecology* 33 (2): 206–17. <https://doi.org/10.1111/1365-2435.13237>.

- Kalluri, Pratyusha. 2020. "Don't ask if artificial intelligence is good or fair, ask how it shifts power." *Nature* 583 (7815): 169–69. <https://doi.org/10.1038/d41586-020-02003-2>.
- Kivelä, Mikko, Alex Arenas, Marc Barthelemy, James P Gleeson, Yamir Moreno, and Mason A Porter. 2014. "Multilayer networks." *Journal of complex networks* 2 (3): 203–71. <https://doi.org/10.1093/comnet/cnu016>.
- Magelinski, Thomas, Mihovil Bartulovic, and Kathleen M Carley. 2021. "Measuring Node Contribution to Community Structure With Modularity Vitality." *IEEE Transactions on Network Science and Engineering* 8 (1): 707–23. <https://doi.org/10.1109/TNSE.2020.3049068>.
- Magelinski, Thomas, David M Beskow, and Kathleen M Carley. 2020. "Graph-Hist: Graph Classification from Latent Feature Histograms with Application to Bot Detection." In *AAAI*, 5134–41. <https://doi.org/10.1609/aaai.v34i04.5956>.
- Magelinski, Thomas, and Kathleen M Carley. 2020. *Detecting Coordinated Behavior in the Twitter Campaign to Reopen America*. Social-Cybersecurity in Times of Crisis and Change.
- Mateescu, Alexandra, Douglas Brunton, Alex Rosenblat, Desmond Patton, Zachary Gold, and danah boyd danah. 2015. "Social media surveillance and law enforcement." *Data Civ Rights* 27:2015–27. [https://datasociety.net/wp-content/uploads/2015/10/Social\\_Media\\_Surveillance\\_and\\_Law\\_Enforcement.pdf](https://datasociety.net/wp-content/uploads/2015/10/Social_Media_Surveillance_and_Law_Enforcement.pdf).
- Ng, Lynnette Hui Xian, Iain Cruickshank, and Kathleen M Carley. 2021. "Coordinating Narratives and the Capitol Riots on Parler." *arXiv preprint*, arXiv: 2109.00945.
- Nizzoli, Leonardo, Serena Tardelli, Marco Avvenuti, Stefano Cresci, and Maurizio Tesconi. 2020. "Coordinated behavior on social media in 2019 UK general election." *arXiv preprint*, arXiv: 2008.08370.
- Owen, Stephen. 2017. "Monitoring social media and protest movements: Ensuring political order through surveillance and surveillance discourse." *Social Identities* 23 (6): 688–700. <https://doi.org/10.1080/13504630.2017.1291092>.
- Pacheco, Diogo, Alessandro Flammini, and Filippo Menczer. 2020. "Unveiling coordinated groups behind white helmets disinformation." In *Companion Proceedings of the Web Conference 2020*, 611–16. <https://doi.org/10.1145/3366424.3385775>.
- Pacheco, Diogo, Pik-Mai Hui, Christopher Torres-Lugo, Bao Tran Truong, Alessandro Flammini, and Filippo Menczer. 2020. "Uncovering Coordinated Networks on Social Media." *arXiv preprint*, arXiv: 2001.05658.
- Patton, Desmond Upton, Douglas-Wade Brunton, Andrea Dixon, Reuben Jonathan Miller, Patrick Leonard, and Rose Hackman. 2017. "Stop and frisk online: theorizing everyday racism in digital policing in the use of social media for identification of criminal conduct and associations." *Social Media+ Society* 3 (3): 1–10. <https://doi.org/10.1177/2056305117733344>.
- Qin, Bei, David Strömberg, and Yanhui Wu. 2017. "Why does China allow freer social media? Protests versus surveillance and propaganda." *Journal of Economic Perspectives* 31 (1): 117–40. <https://doi.org/10.1257/jep.31.1.117>.
- Russell, Stuart. 2019. *Human compatible: Artificial intelligence and the problem of control*. Penguin.

- Shepherd, Katie, and Moriah Balingit. 2020. *A noose, an ax and trump-inspired insults: Anti-lockdown protesters ratchet up violent rhetoric*, May. <https://www.washingtonpost.com/nation/2020/05/15/noose-fight-coronavirus-protest/>.
- Skorniakov, Kirill, Denis Turdakov, and Andrey Zhabotinsky. 2018. "Make Social Networks Clean Again: Graph Embedding and Stacking Classifiers for Bot Detection." In *CIKM Workshops*.
- Stanley-Becker, Isaac, and Tony Romm. 2020. *The anti-quarantine protests seem spontaneous. but behind the scenes, a powerful network is helping.*, April. [https://www.washingtonpost.com/politics/inside-the-conservative-networks-backing-anti-quarantine-protests/2020/04/22/da75c81e-83fe-11ea-a3eb-e9fc93160703\\_story.html](https://www.washingtonpost.com/politics/inside-the-conservative-networks-backing-anti-quarantine-protests/2020/04/22/da75c81e-83fe-11ea-a3eb-e9fc93160703_story.html).
- Steinert-Threlkeld, Zachary C, Delia Mocanu, Alessandro Vespignani, and James Fowler. 2015. "Online social networks and offline protest." *EPJ Data Science* 4 (1): 1–9. <https://doi.org/10.1140/epjds/s13688-015-0056-y>.
- Swain, Val. 2013. "Disruption policing: surveillance and the right to protest." *Open Democracy* 8 (08): 2013. <https://www.opendemocracy.net/en/opensecurity/disruption-policing-surveillance-and-right-to-protest/>.
- Tracy, Marc. 2020. *Anti-lockdown protesters get in reporters' (masked) faces*, May. <https://www.nytimes.com/2020/05/13/business/media/lockdown-protests-reporters.html>.
- Vargas, Luis, Patrick Emami, and Patrick Traynor. 2020. "On the detection of disinformation campaign activity with network analysis." In *Proceedings of the 2020 ACM SIGSAC Conference on Cloud Computing Security Workshop*, 133–46. <https://doi.org/10.1145/3411495.3421363>. arXiv: 2005.13466v2.
- Wasserman, Stanley, Katherine Faust, et al. 1994. *Social network analysis: Methods and applications*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511815478>.
- Weber, Derek, and Lucia Falzon. 2021. "Temporal nuances of coordination networks." *arXiv preprint arXiv:2107.02588*.
- Weber, Derek, and Frank Neumann. 2020. "Who's in the gang? Revealing coordinating communities in social media." In *2020 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, 89–93. IEEE. arXiv: 2010.08180v1.
- Weiss, Brennan. 2018. *Russian trolls orchestrated divisive protests in the US about Trump - here are 9 that we know about*, February. <https://www.businessinsider.com/russians-organized-pro-anti-trump-rallies-to-sow-discord-2018-2#july-23-2016-new-york-3>.
- Williams, Matthew L, Adam Edwards, William Housley, Peter Burnap, Omer Rana, Nick Avis, Jeffrey Morgan, and Luke Sloan. 2013. "Policing cyber-neighbourhoods: tension monitoring and social media networks." *Policing and society* 23 (4): 461–81. <https://doi.org/10.1080/10439463.2013.780225>.
- Yang, Kai-Cheng, Onur Varol, Pik-Mai Hui, and Filippo Menczer. 2020. "Scalable and generalizable social bot detection through data selection." In *Proceedings of the AAAI Conference on Artificial Intelligence*, 34:1096–103. 01. <https://doi.org/10.1609/aaai.v34i01.5460>.

- Yang, Lei, Tao Sun, Ming Zhang, and Qiaozhu Mei. 2012. "We know what@ you# tag: does the dual role affect hashtag adoption?" In *Proceedings of the 21st international conference on World Wide Web*, 261–70. <https://doi.org/10.1145/2187836.2187872>.
- Yang, Yan, and Hao Wang. 2018. "Multi-view clustering: A survey." *Big Data Mining and Analytics* 1 (2): 83–107. <https://doi.org/10.26599/BDMA.2018.9020003>.
- Zhang, Han, Chang-Dong Wang, Jian-Huang Lai, and S Yu Philip. 2017. "Modularity in complex multilayer networks with multiple aspects: a static perspective." In *Applied Informatics*, 4:1–29. 1. SpringerOpen. <https://doi.org/10.1186/s40535-017-0035-4>.

## Authors

**Thomas Magelinski** (B.S. Virginia Tech) is a Societal Computing PhD student at Carnegie Mellon University and is advised by Professor Kathleen Carley. He is a Knight Fellow in the Center for Informed Democracy & Social Cybersecurity (IDeaS), and a graduate researcher in the Computational Analysis of Social and Organizational Systems (CASOS). His research draws on network science and machine learning to build methods for the analysis of dynamic heterogeneous online social networks.

**Lynnette Hui Xian Ng** is a graduate researcher at the Center for Informed Democracy & Social Cybersecurity (IDeaS) and Computational Analysis of Social and Organizational Systems (CASOS), at Carnegie Mellon University. Her research examines social cybersecurity, digital disinformation and online opinion synchronization.

**Kathleen M. Carley** (H.D. University of Zurich, Ph.D. Harvard, S.B. MIT) is a Professor of Societal Computing, Institute for Software Research, Carnegie-Mellon University; Director of the Center for Computational Analysis of Social and Organizational Systems (CASOS), Director of the Center for Informed Democracy and Social Cybersecurity (IDeaS), and CEO of Netanomics. Her research blends computer science and social science to address complex real world issues such as social cybersecurity, disinformation, disease contagion, disaster response, and terrorism from a high dimensional network analytic, machine learning, and natural language processing perspective. She and her groups have developed network and simulation tools, such as ORA, that can assess network and social media data.

## Acknowledgements

We would like to thank Matthew Babcock, PhD for collecting the Reopen America Twitter dataset used for this study.

## Data Availability Statement

The collected dataset is not made publicly available due to the Twitter API's terms of service. However the code to replicate results on other Twitter datasets is available at <https://github.com/CASOS-IDeaS-CMU/coordination-analysis>.

## Funding Statement

This work was supported in part by the Office of Naval Research (ONR) Award N000141512797 Minerva award for Dynamic Statistical Network Informatics, and the Center for Computational Analysis of Social and Organization Systems (CASOS). Thomas Magelinski was also supported by the IDeaS Center as a Knight Fellow. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the ONR.

## Ethical Standards

We discussed the problem of coordination detection on social media and its ethical implications in Section 3, together with the need for responsible coordination detectors. During data collection, we only collected tweets that were made public on Twitter. In this work, we masked author names except for news and highly influential accounts.

**Keywords**

coordination detection, social media, synchronized action framework