



# The coordination network toolkit: a framework for detecting and analysing coordinated behaviour on social media

Timothy Graham<sup>1</sup> · Sam Hames<sup>2</sup> · Elizabeth Alpert<sup>3</sup>

Received: 5 November 2023 / Accepted: 22 February 2024  
© The Author(s) 2024

## Abstract

This paper introduces and evaluates the Coordination Network Toolkit, an open-source software package and methodological framework designed to detect and analyse coordinated behaviour on social media platforms. As the dynamics of online communication continue to evolve, coordination analysis has emerged as an important field of study with significant implications for understanding online influence, digital astroturfing, and online activism. Recognising the absence of a comprehensive, open-source tool for constructing coordination networks, our approach fills this gap, catering to multiple behaviors across diverse social media platforms. Our approach synthesises and significantly enhances various methods to provide a methodological framework for ‘multi-behaviour’ coordination detection, utilising weighted, directed multigraphs to capture intricate coordination dynamics. We evaluate our approach by revisiting a case study of the 2020 #ReopenAmerica Covid protest movement on Twitter. The paper concludes with a set of recommendations for future work, emphasising the need for a tailored statistical framework for coordination analysis and a deeper exploration into the motives behind online coordination.

**Keywords** Coordinated behaviour · Coordination · Social media · Disinformation · Astroturfing · Coordinated inauthentic behaviour

## Introduction

By design, online influence campaigns are difficult to discover and analyse. Whilst still a relatively new field, researchers draw on an increasingly standardised suite of qualitative and open-source intelligence methods and tools to discover and analyse suspected influence campaigns. Among these, social network analysis (SNA) has emerged as a popular and highly effective approach to mixed methods forensic analysis of information operations. By collecting large-scale

---

Extended author information available on the last page of the article

digital datasets, practitioners can examine the structure of interactions between accounts involved in influence campaigns on platforms such as Twitter in order to map and understand their scale and scope. For instance, retweet networks were popularly used during the 2016 US presidential election to analyse Russian state-backed troll activity [1, 2] and have since become a de facto feature of quantitative approaches to studying online behaviour.

A recent innovation in SNA has been the study of coordinated behaviour on social media. Various referred to as coordination networks [3], latent coordination networks [4], coordinated networks [5, 6] and synchronised action networks [7], this approach involves mapping similar activities performed by groups of accounts repeatedly within a short time window of one another. For example, co-retweet networks provide insight into coordinated content amplification on Twitter, where Twitter accounts (nodes) in the network are connected by an link or edge if they exhibit a pattern of retweeting the same tweet within 60 s of one another [8, 9]. This methodological approach appears to be broadly effective at revealing online influence campaigns such as political astroturfing [10], where fake accounts attempt to hijack online conversations and simulate grassroots support for a political party, person, ideology, idea, or event.

The majority of approaches to detecting and analysing coordination focus on single behaviours at a time on single platforms. For example, co-retweet networks map coordinated retweeting on X (formerly Twitter) whereby accounts repost the same tweet repeatedly within a short time window of each other [8]. Similarly, co-tweet networks map coordinated tweeting whereby accounts post identical messages repeatedly within a short time window [10]. Co-link networks map coordinated link sharing on Facebook where accounts post the same links repeatedly within a short time window [11]. The case studies in [6] examine single behaviours in their distinct contexts. The authors highlight this limitation, suggesting that ‘more work can be done in considering multiple dimensions of coordination in specific scenarios’ [6, p. 465]. There are primarily two approaches in the literature that afford mapping multiple coordinated behaviours on social media.

Firstly, the study by [7] propose handling multiple behaviours through the ‘synchronised action framework’. They propose the use of multi-level networks to map different action types, where an action is any measurable behaviour that can be associated with a particular time point. The authors focus on emergent behaviours that can be derived from the post text. The edges in the Multi-View Coordination Networks approach are undirected and weighted. Secondly, the proposed approach by [4] addresses multiple behaviour types through a method known as ‘latent coordination networks’ (LCN). In this approach, LCNs are built by mapping different ‘interaction types’ between accounts in a social media dataset. Interaction types include those observed in the literature such as co-retweet, co-tweet, co-link, but also co-mentions (mentioning the same account ID) and co-conv (joining the same reply thread). The authors note that for many community detection tools to be effective, the LCN might need simplification, wherein multiple edges with different interaction types are reduced to single weighted edges. Following this procedure, [4] identify ‘highly coordinating communities’ (HCC), labeled as  $H$ , within the LCN,  $L$ .

A suitable community detection algorithm is employed for this purpose. Both the HCCs and LCNs are undirected, weighted graphs.

### Limitations of current approaches

A general limitation of existing approaches concerns how relationships are mapped as edges in the networks. In previous approaches, edges are undirected and weighted, meaning that we only know that a given pair of coordinating nodes  $x$  and  $y$  have coordinated  $n$  times for action-type  $i$  within a time window  $t$ . Clearly this poses a ‘time window problem’ for pairs of nodes where, for example, one is spammy and performs many actions within  $t$  and the other only performs a couple of actions within the same time window. As [7] state: this means “spammers will be erroneously strongly connected to normal users” (p. 7). One solution offered by [7] is to only connect a node to other nodes if it has a lower presence in the time window than the other potential connecting nodes. For example, if account  $x$  has tweeted twice in a time window  $t$  and the other user  $y$  has tweeted 100 times, it would be erroneous to create a strongly weighted edge between  $x$  and  $y$  unless  $y$  has tweeted again later in the window. Otherwise  $y$  will have strongly weighted connections with all eligible users who happened to tweet with the same behaviour (or action-type) within  $t$ —they are ‘caught in the net’ of  $y$ ’s spammy behaviour. However, this solution leads to information loss as it is a consequence of edge window effects. Simply not making connections between nodes based on the order of posting results in information loss as such asymmetric posting behaviours are of analytical importance. For example, a bot net might always retweet a target account to amplify its content artificially—in this case, the bots will always post after the target and this asymmetrical relationship is important to retain in the network. Similarly, a target might be the subject of coordinated abuse or ‘brigading’ [12] where coordinated accounts spam replies on every post the target makes—again, this is critical to map in the network and should not be omitted.

A second consideration for previous work are the limitations for handling multiple coordinating behaviours simultaneously. While multi-level networks proposed by [7] offer some advantages for clustering complex or higher-order behaviours in social media data, they are also specialised types of networks that require specific functionality and algorithms to compute over the entire graph. Switching between the levels or hierarchies of the behaviour types complicates the task of assessing what kinds of coordinated behaviour are occurring simultaneously within and between clusters, as discussed previously. A similar approach using two-level networks is suggested by [4] who propose ‘account-reason’ networks to model multiple behaviour types. In the account-reason network, nodes have two types (accounts and ‘reasons’ for linking). Edges also have two types, denoting either ‘coordination’ (edge type 1) between account nodes or ‘associated because’ (edge type 2) between account nodes and reason nodes. For example, reason nodes might be ‘co-retweet and co-link’ or ‘co-conv and co-mention’ or a single reason such as ‘co-tweet’, or any combination depending on the network construction choices. The analysis of such multi-level networks is left to future work and the

authors simply note that deeper insights may be revealed using multi-level network analysis techniques [4], p. 465].

The multilevel graph structure proposed by [7] and [4] makes it difficult to compare the results with existing methods in the field, which use a more standard graph structure and individual behaviour types in separate graphs. This produces a problem of commensurability whereby comparing the results becomes an ‘apples and oranges’ task due to differences in network construction and analytical procedure. While there is no ‘standard’ for coordinated behaviour analysis, the Multi-View Coordination approach is sufficiently different from other approaches outlined above that it makes working with such networks more difficult due to the specialised algorithms and software required and complicates the task of comparing results with previous studies. Separating behaviour or ‘action’ types into their own layers means that specialised clustering techniques [13] must be used to collapse the layers into a single network consisting of clusters of interest due to their centrality scores and/or density properties. This poses a problem for analysts because it results in information loss—what we are looking at is no longer the coordinated behaviour but a summarised clustered view of it. Indeed, [7] find little evidence of different behaviours (i.e., ‘action-types’) within coordinating clusters derived from collapsing the layers using the multi-view modularity approach. In real-world case studies it is possible—likely even—that inauthentic accounts in orchestrated campaigns are simultaneously coordinating in multiple ways.

## **The need for a synthesised, open-source coordination network toolkit**

The toolkit and methodological framework we develop in this paper addresses several key gaps in the field. While there are various existing approaches to constructing coordination networks, they tend to either be focused on single behaviours and/or single platforms. While at least two approaches allow for multiple behaviours, these rely on clustering and multi-level network approaches that do not afford analysis of multiple behaviour simultaneously in the same view. Few software implementations exist and those that do focus on single behaviours and platforms.

Against this backdrop, we identify the need for an open-source toolkit and methodological framework that synthesises the various approaches and offers a standardised, computationally efficient software toolkit. This toolkit seeks to complement and significantly extend the state-of-the-art in the field by making it ‘backwards compatible’ with existing approaches while also offering improvements in several key ways. The paper is thereby guided by the following questions:

- RQ1 How can we implement an approach to constructing coordination networks that is flexible and optimised for large-scale datasets?
- RQ2 How can we address the time window problem of ‘spammy’ accounts in coordination networks?

**RQ3** How can we develop a network structure (i.e., a type of graph) that maps multiple behaviours simultaneously, which can be visualised and analysed using standard open-source software for network analysis?

## A standardised framework for coordination network analysis

The Coordination Network Toolkit<sup>1</sup> includes functionality for co-tweeting and co-retweeting [8, 9], where accounts post exactly the same text (co-tweets) or repost the same post within a short time window (co-retweets). Secondly, it includes functionality for co-link analysis, where multiple accounts post the same URLs repeatedly and in a short time window of each other [5, 11]. Thirdly, it adds three new types of network types: co-reply, where accounts are replying to the same post repeatedly together; co-similarity, where accounts post similar text (but not exact duplicates), which relaxes the strict assumption of co-tweeting; and co-posting, where accounts make any post within the same time window. Co-posting is not likely to be directly useful for identifying coordination but does provide insight into the maximum possible coordination strengths between nodes in the network based on the temporal dynamics of posting.

### Computational approach

To solve the time-window problem discussed in Sect. 1.1, the Coordination Network Toolkit implements directed and weighted edges for coordinating accounts using a sliding time window approach. This approach avoids window edge effects caused by examining coordination in fixed global time windows by looking for coordination in a window of time centered on the specific message of interest. Additionally, it handles the case of coordinating events happening at different rates within that time window by constructing an asymmetric network. Consider two accounts  $x$  and  $y$ : if  $x$  posts 100 times within a time window, and  $y$  posts only once, we cannot consider this as an undirected network because of the extreme asymmetry in volume. Therefore we construct the network as a directed network where  $x$  is considered to have coordinated with  $y$  100 times in that time window and therefore has a 100 weight edge from  $x$  to  $y$ , but  $y$  has only coordinated with  $x$  once with a corresponding reverse edge of weight 1. Edge weights are therefore asymmetric using directed edges, to allow detailed examination and handling of uneven posting behaviour—this is of interest to analysts and therefore such information should be conserved in the network. This is particularly important for considering abusive forms of coordination: an account  $z$  subject to a coordinated attack of automated replies should not be said to be coordinating (have outbound edges) with the accounts replying to them.

<sup>1</sup> <https://github.com/QUT-Digital-Observatory/coordination-network-toolkit>

A simple computation over a social media dataset to detect coordinating pairs of accounts would require  $O(n^2)$  account comparisons. Efficient computation therefore requires an awareness of the data locality of the different operations and selection of effective ordering of data for locality of reference. Fortunately existing relational databases are already capable of doing this in a flexible way by constructing indexes of appropriate structure for the given calculation. Text equality comparisons (for co-tweet) are also accelerated by comparing a checksum first, before checking the actual text. This eliminates text that can't match early and leads to a compact index for computational purposes. The approach used to construct these networks then amounts to:

- Choosing an appropriate data locality for the specific problem.
- Constructing an appropriate index in the database to make it possible to use that locality in computation.
- Enabling parallelisation—the size of intermediate results is controlled by computing outbound edges node-wise. This allows for a flexible, extensible framework for computing and aggregating partial results in parallel, while using minimal memory.

The toolkit is implemented in Python and SQLite, with parallel processing for performance. SQLite is used as a persistent data store for three reasons: (1) because of its fundamental affordances for custom indexing of tables to support the required sort order for different calculations (2) it is a reliable persistent data store, meaning that preprocessing and other data preparation steps such as deduplication only need to be run once and can be reused to construct many different networks (3) As a relational database originally intended for embedded use cases it can be very light on memory usage. The Python layer on top of SQLite enables parallel processing by breaking down the network calculation into operations on the nodes (accounts), and computing outbound edges for batches of nodes in parallel. The Python layer also allows for customisation of certain features, with a particular focus on ensuring that text processing can be adapted to suit different needs for co-tweet and similarity calculations.

The open-source toolkit support platforms that are organised around user accounts posting messages with distinct identifiers. Most testing is done with X (formerly Twitter) and X API data is supported natively. General CSV support is available for any platform. Users supply data in the appropriate CSV format and the toolkit can compute over it to construct coordination networks. The CSV format is intended to support many common affordances for social media platforms such as replies to other messages or inclusion of URLs. To use the CSV ingest format, users can construct a CSV with a header and the following columns. The names of the columns don't matter but the order does.

- `message_id`: the unique identifier of the message on the platform.
- `user_id`: the unique identifier of the user on the platform.
- `username`: the text of the username (only used for display).

- `repost_id`: if the message is a verbatim report of another message (such as a retweet or reblog), this is the identifier of that other message. Empty strings will be converted to null.
- `reply_id`: if the message is in reply to another message, the identifier for that other message. Empty strings will be converted to null.
- `message`: the text of the message.
- `timestamp`: A timestamp in seconds for the message. The absolute offset does not matter, but it needs to be consistent across all rows.
- `urls`: A space delimited string containing all of the URLs in the message

### Handling text: co-tweet and co-similarity

An obvious approach to identify co-tweeting behaviour [10] is to perform a simple string comparison for equality. This works for the most extreme cases of inauthentic coordinated behaviour where accounts are posting exact duplicate text, but can fail when even minor alterations occur. These alterations can include use of basic affordances of the platform, such as using @ mentions to message a particular user. To handle these common cases we extend the co-tweet idea further to handle some common transformations of the the text: by default we perform co-tweet analysis by comparing strings after lower casing, removing of @mentions and standardisation of whitespace—this ensures that we capture common behaviours such as copying and pasting the same message, but directing them at different users, or replying to multiple threads with the an otherwise identical message.

We also extend the idea of co-tweeting to search for messages with highly similar but not identical content: in this mode tweets are tokenised, and the Jaccard similarity between the sets of tokens in two messages within the target time window is calculated. While the current implementation in the toolkit is basic, this approach could be extended to use more sophisticated measures of similarity using the toolkit's inbuilt customisation hooks.

### Performance

This implementation is highly scalable and can make use of most available CPU resources on standard desktops and laptops: it is also memory efficient. Exact runtimes depend on the data of interest and the desired network settings: wider time windows are necessarily going to take longer to compute. While we have not conducted extensive performance comparison we can provide the following indicative timing of certain operations to pre-process and compute various networks on 7.2 million tweets from the timelines of 0.51 million users collected over a 24 h period (see Table 1). All calculations were performed on a standard laptop with a mobile processor (Intel i7 1260p), with all available CPUs allowed to be used.

**Table 1** Processing Times for Calculating Coordination Networks

Procedure	Time Taken
One-off pre-processing	10m54.501s
Co-link network, 300 s time window	0m17.7s
Co-tweet network, 300 s time window	3m4.6s
Co-reply network, 300 s time window	0m27.2s

## Multi-behaviour networks

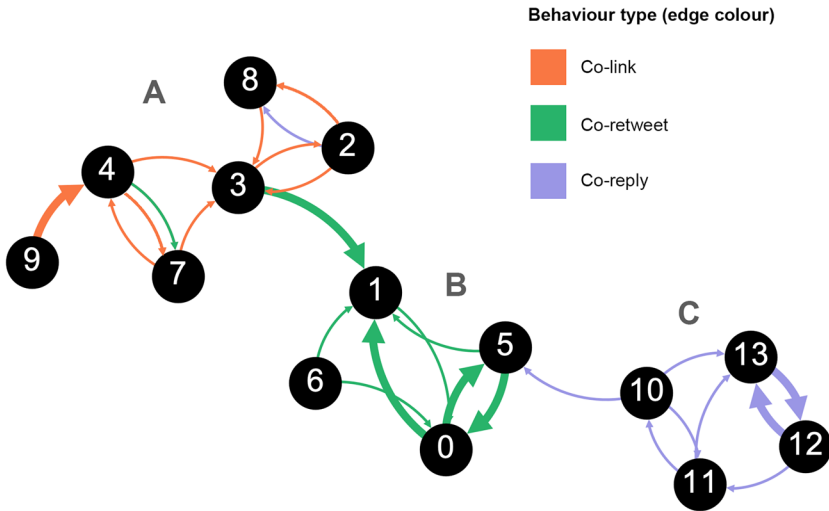
One of the main advantages of the Coordination Network Toolkit is the ability to construct coordination networks that include multiple types of behaviour simultaneously. To achieve this, we propose labelled multigraphs with directed and weighted edges (herein labelled multidigraph) to construct coordination networks that map multiple behaviour types simultaneously. We call this the Multi-behaviour Coordination Network.<sup>2</sup>

A labelled multidigraph allows node pairs  $(u, v)$  to have multiple edges with the same direction between them and a map that describes the label of each edge from a finite alphabet. The application of labelled multidigraphs to coordination networks is novel because it allows a given ordered pair of coordinating vertices  $(u, v)$  to have multiple behaviour types simultaneously, where each behaviour type is designated by a label on a directed and weighted edge representing the number of times  $u$  coordinated with  $v$  in a sliding time window  $t$  (where  $t$  is usually  $t = 60$  seconds). In this way, a Multi-behaviour Coordination Network consists of a set of vertices  $V = \{v_1, v_2, v_3, \dots, v_n\}$ , edges  $E = \{e_1, e_2, e_3, \dots, e_n\}$ , edge weights  $w : E \rightarrow \mathbb{Z}^+$  and a function  $f : E \rightarrow V \times V = \{(u, v) | u, v \in V\}$  such that the edges  $e_1$  and  $e_2$  are multiple edges if  $f(e_1) = f(e_2)$ , that is, the head and tail vertices of  $e_1$  and  $e_2$  are identical. The vertices are user accounts and edges are the directed and weighted coordination links between accounts, calculated using the methods described previously and outlined further below. The nodes and edges are labelled from two finite alphabets  $\Sigma_V$  and  $\Sigma_E$  that provide the available labels for the accounts (i.e. account IDs or usernames) and edges (i.e., behaviour type) respectively. The labelling of the nodes and edges are described by two maps  $\Phi_V \rightarrow \Sigma_V$  and  $\Phi_E \rightarrow \Sigma_E$ . Hence a Multi-behaviour Coordination Network is a 7-tuple  $G = \{V, E, w, \Sigma_V, \Sigma_E, \Phi_V, \Phi_E\}$ .

Figure 1 shows an example of a Multi-behaviour Coordination Network visualised using the Gephi software [14] using the Force Atlas 2 layout algorithm [15]. The vertices  $V$  are Twitter accounts labelled  $\Sigma_V = \{0, 1, \dots, 12\}$  and weighted edges  $E$  are labelled by behaviour type  $\Sigma_E = \{\text{co-link, co-retweet, co-reply}\}$  and coloured accordingly in the set  $\{\text{red, green, blue}\}$ . This example network shows three behaviour types. The first is coordinated link sharing or co-linking (red edges). The second

<sup>2</sup> At the time of writing, the open-source toolkit does not include functionality to generate Multi-behaviour Coordination Networks. To achieve this, we use the networkx package in Python to combine the individual coordination network types and use edge labels to denote behaviour type, as outlined in this section.





**Fig. 1** An example of a Multi-behaviour Coordination Network. There are three types of behaviours shown by the edge colouring: co-linking (red), co-retweeting (green), and co-replying (blue)

is coordinated retweeting or co-retweeting (green edges). The third behaviour type is coordinated replying (blue edges). In this example we observe the behaviours structured naturally into three clusters, annotated A, B, and C in Fig. 1. However, we also observe that some nodes are engaged in multiple behaviours. For example, in Cluster A node 4 and node 7 are engaged in coordinated link sharing (co-linking) behaviour, as shown by the two red-coloured edges directed from 4 to 7 and directed from 7 to 4. However, 4 is also asymmetrically coordinating with 7 for coordinated retweeting behaviour (co-retweeting) as shown by the green edge directed from 4 to 7. It is asymmetric coordination because 7 does not reciprocate with 4 for co-retweeting behaviour and therefore no edge is directed from 7 to 4. Similarly, in Cluster A node 2 is asymmetrically coordinating with 8 for both co-link and co-reply behaviour—there are two directed edges going from 2 to 8.

Multi-behaviour Coordination Network shown in Fig. 1 maps multiple behaviours in a single view and is fully compatible with standard graph algorithms that support directed, weighted edges. This reduces the information loss associated with the necessity of clustering multi-level networks into ‘flat’ networks, as per [7]. It also addresses the limitation of other approaches which use undirected weighted edges to map coordinated behaviour. The directionality of edges in the Multi-behaviour Coordination Network provides information about asymmetric posting behaviour which is useful for forensic analysis of networks and algorithmically for computation. For example, in Fig. 1 the asymmetric co-retweeting behaviour between node 4 and node 7 may suggest that node 4 is using automation or tweet scheduling software to automatically retweet anything that node 7 retweets. The directionality of the edge provides important information about the content production processes in Cluster A, which is otherwise lost in undirected networks. Likewise, the Multi-behaviour Coordination Network approach is in a sense ‘backwards compatible’

with key concepts and existing methods in the literature on coordination networks. While there are some important differences with edge construction, datasets analysed with the Multi-behaviour Coordination Network approach can be compared with single-behaviour networks such as co-link networks [11] or co-retweet networks [8], given that each type of behaviour is a subgraph filtered by edge type. For example, to retrieve a co-retweet network from a Multi-behaviour Coordination Network, the user simply filters by edge type = co-retweet and removes nodes with 0 degree (i.e., nodes that no longer have any edges because they were removed).

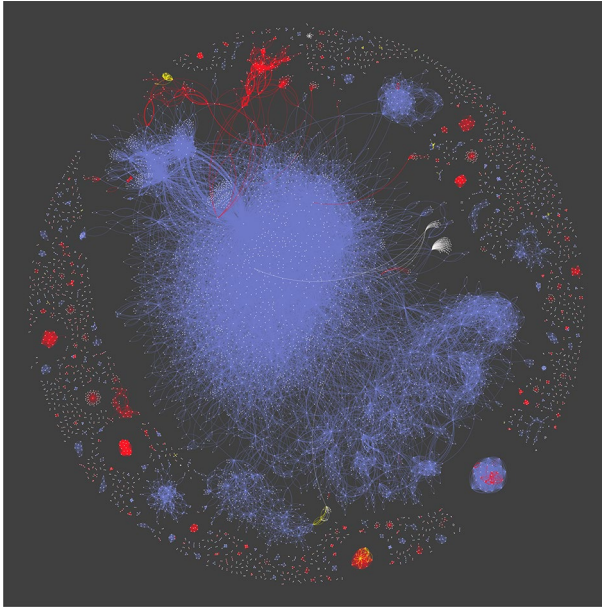
## Case study of #ReopenAmerica protests

In this section, we undertake a case study analysis of coordinated activity on Twitter during the “Reopen America” protests. This study reproduces the analysis of [7] who used this event as a way to demonstrate the effectiveness of their approach, although we use more standardised types of behaviour as detailed below. The Reopen America protests took place from April to September 2020. This protest movement sought to ease the COVID-19 safety restrictions that impacted daily life and activities across the United States. The demonstrations were marked by some participants neglecting safety protocols, such as social distancing and mask-wearing, and confrontations with media personnel. Notably, the consistency in organizing materials, event dates, and phrasing across states sparked inquiries about potential coordinated efforts. [7] suggest that this makes the event interesting as a case study for analysing coordinated activity on social media—in this case Twitter.

The authors of the original study leveraged the Twitter Search API to amass a comprehensive dataset of 9.9 millions tweets, spanning from April 1 to June 22, 2020. The tweets were authored by 3.6 million unique accounts. Their search strategy incorporated keywords and hashtags such as ‘open up’, ‘reopen’, ‘operation gridlock’, and ‘liberate’. To ensure a state-specific analysis, they appended all US state abbreviations to these terms (e.g., ‘liberateNY’). These terms were derived via preliminary observations of trending hashtags related to the protests. While the methodology primarily aimed to encapsulate tweets discussing the protests, it inevitably encompassed other related discussions, a recognised limitation of keyword-based data collection. Despite this, the authors identified coordinated activity across various discussions, both directly linked to the protests and otherwise.

In our current study, we gathered a Twitter dataset during the same period and using an identical set of keywords. This resulted in a dataset of 9,058,611 tweets sent by 2,949,118 unique accounts. While we are not reproducing the exact methods devised by [7], we use this case study as an opportunity to compare the results of Multi-behaviour Coordination Network analysis using the open-source toolkit we have developed. By comparing our findings with the original study, we aspire to both validate the initial observations and underscore the efficacy and contribution of our approach.

Figure 2 shows a force-directed network visualisation of coordinated activity around the #ReopenAmerica protest movement. The nodes are pairs of accounts in the dataset who coordinated at least five times within 300 s of each other, which

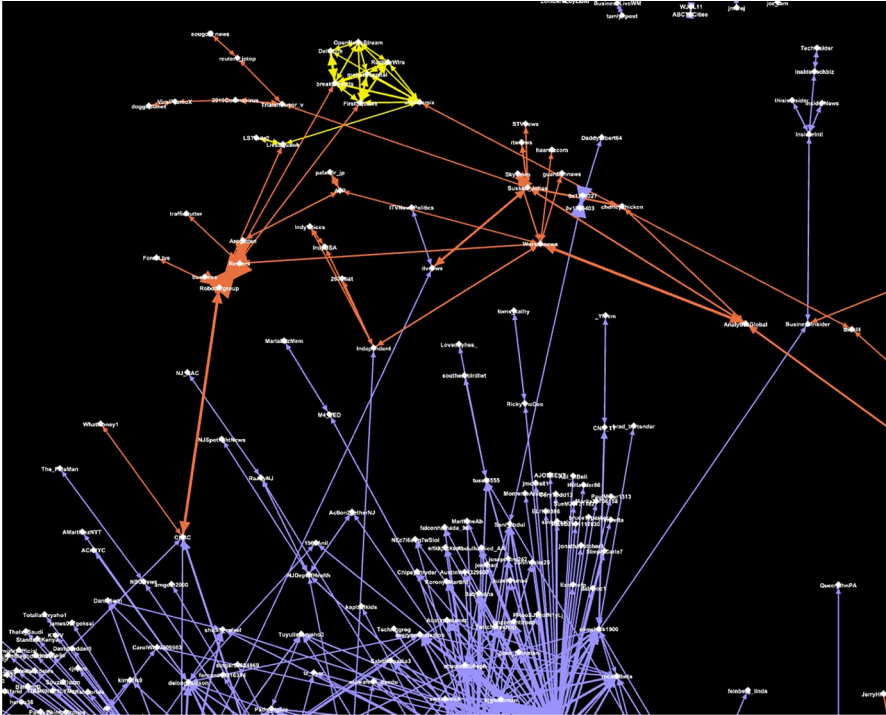


**Fig. 2** A network visualisation of the #ReopenAmerica dataset. There are four types of behaviours shown by the edge colouring: co-linking (red), co-retweeting (blue), co-tweeting (yellow) and co-replying (white)

is the same time window employed in the original study. The network contains 7280 nodes connected by 26513 edges. There are four types of coordination fully observed in this network: coordinated link sharing or co-linking (red); coordinated retweeting or co-retweeting (blue), coordinated tweeting or co-tweeting where accounts send duplicated messages (yellow), and coordinated replying or co-replying where accounts reply to the same tweet (white).

Following [7] we identify three communities of interest in this network. We use eigenvector centrality and average clustering coefficient to identify clusters that are both highly modular and also central in the network. To this we add a further feature of interest: multiple types of coordinated behaviour in the same cluster. Rather than having different ‘views’ of the data that analysts must explore individually, our approach provides the ability to analyse all the observed coordinated activity in the network in a single view. This is a marked improvement over previous methods as our approach enables users to visually and analytically disambiguate multiple behaviour types within the same network view, as Fig. 2 clearly demonstrates.

First we turn attention to the giant connected component in the middle of Fig. 2. This is a co-retweet network as shown by the blue edge colouring denoting the behaviour type. Interestingly, this is a counter-public to the #ReopenAmerica protest movement - we might label this cluster ‘COVID safety advocates’ or even ‘Reopen responsibly advocates’. The tweets predominantly reflect concerns about hasty reopening, emphasise the importance of safety protocols, and critique political decisions that might endanger public health. The discourse in this cluster promotes a



**Fig. 3** A subnetwork of the #ReopenAmerica dataset showing three types of coordinated behaviour within the same view: co-linking (red), co-retweeting (blue), and co-tweeting (yellow)

more cautious, data-driven, and safety-first approach to reopening, in contrast to the more aggressive push for rapid reopening seen in the protests. The network contains the nodes with the highest eigenvector centrality, although it is not as modular as many of the smaller clusters that we discuss further below. However, what is particularly interesting about this giant connected component of co-retweeting accounts is how it connects into the rest of the network. Visually we can see from Fig. 2 how this cluster has tendrils of connections that reach out into other areas of the network. In particular, we are interested in how it connects to clusters where other behaviour types are occurring (Fig. 3).

Figure 4 shows a subnetwork attached to the ‘COVID safety advocates’ co-retweet public—this is a zoomed-in portion of the top-part of Fig. 2. This subnetwork is interesting because it shows coordination of three varieties: co-linking (red), co-retweeting (blue), and co-tweeting (yellow). We can visually trace the connections between accounts to understand the scale and scope of coordinated activity across multiple behaviour types. We discover that two major newspapers—CNBC and The Independent—are the connecting points from the co-retweet counter-public to this subnetwork. The coordinated link network (red edges) consist of various international and domestic news organisations that repeatedly share links to each others’ content. In turn, the co-link cluster connects into a coordinated tweet or co-tweet cluster (yellow edges) consisting of business and financial news





and co-linking. Based on weighted degree, the broadcasters have engaged in 843 instances of co-linking (17.16% of total) and 4,070 co-retweets (82.84% of total). Structurally, we observe that although only a subset of broadcasters engage in co-link behaviour, there is an embedded clique of co-linkers at the centre of this sub-network. This suggests that there are two strategies of coordinated amplification of Global Television Network content on Twitter and that the multidigraph network structure reveals these organisational logics simultaneously. Compared to previous methods, this enhances the network analysis capabilities because our approach is able to disambiguate messaging strategies within a single view of the network.

Our Multi-behaviour Coordination approach also detects the three campaigns identified by [7].

## 7News Network

In the original paper, [7], Section 5.2.1] performed a visual network analysis and discover that the best-connected component is the Australian 7News network. These 15 accounts tweet the same or nearly identical tweet at the same time to boost their articles to a broader geographic audience across Australia and internationally. This also an additional account that is a different cable news channel known as Sky News Australia (@SkyNewsAust), which used the same hashtags as the 7News accounts. The 7News network covered the ‘reopen’ events in the US with a focus on the eventual reopening of Australia’s borders. [7] found that the 7News network coordinated using three action types: hashtags, URLs, and @mentions. Although we do not consider hashtags and @mentions in our analysis, the Multi-Behaviour network discovers the same network due to the co-linking behaviour of these accounts. This demonstrates that our method captures similar coordination patterns but also adds further insights as we outline below.

Figure 5 shows the 7News coordinated link network, which has an average weighted degree of 229 and average clustering coefficient of 0.971, meaning that it is a locally dense network with 202 out of 210 possible edges formed. Here we observe that the directionality of edges is important because not all the nodes are coordinating in the same way. To explore this further, as shown in Fig. 6 we generated a heatmap to visualise the asymmetric coordination patterns among 7News accounts. The heatmap is structured as a matrix where each cell at position  $(i, j)$  represents the weighted degree difference between account  $i$  and account  $j$ . The value in each cell is computed by taking the weighted out-degree of  $i$  and subtracting the weighted in-degree of  $j$ . In this way, it captures the frequency and intensity of the coordination behavior between pairs of 7News accounts. Red hues in the heatmap signify positive weighted degree differences, emphasizing accounts that post more than others within the time window. Blue hues indicate negative weighted degree differences, highlighting accounts that tend to post less within the same time window as others. The deeper the hue, the stronger the asymmetric coordination behavior between pairs of 7News accounts: a deep red indicates many more outbound messages than inbound messages posted within a

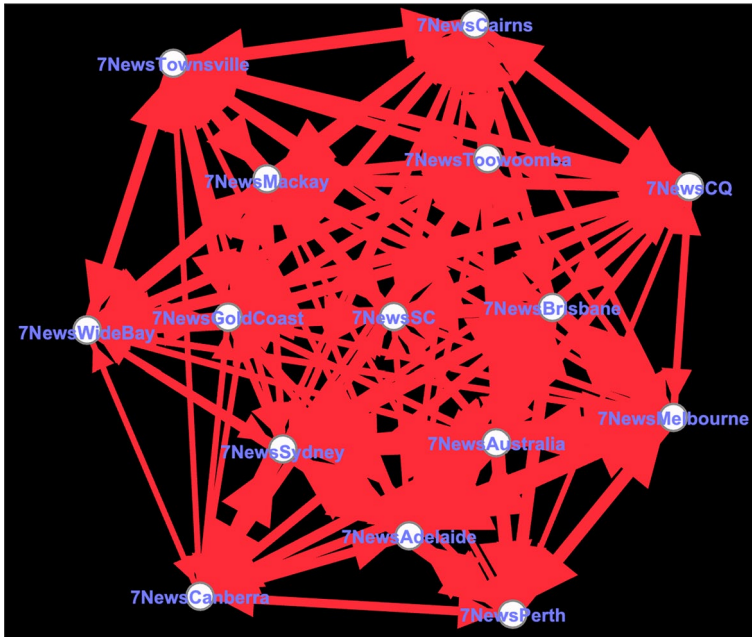
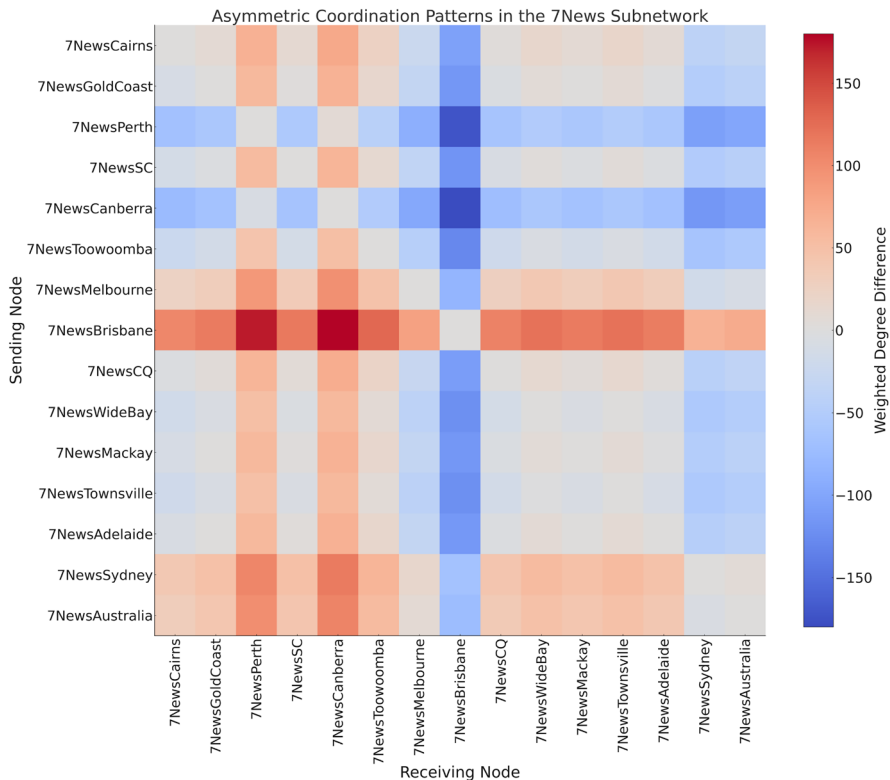


Fig. 5 The 7News Network of Australian news accounts, engaged in a coordinated link-sharing campaign

given time window, a deep blue indicates more inbound messages than outbound messages posted within the time window.

Providing more insight than the methods used by [7], our exploration of the 7News network uncovers pronounced *asymmetric* coordination dynamics, particularly around the ‘7NewsBrisbane’ account. It often posts more than the ‘7NewsCanberra’, ‘7NewsPerth’, ‘7NewsTownsville’, and ‘7NewsWideBay’ accounts within the same time window. Conversely, 7NewsBrisbane tends to follow or synchronize with content from ‘7NewsSydney’, ‘7NewsAustralia’, ‘7NewsMelbourne’, and ‘7NewsCairns’. This indicates a pattern of behaviour where ‘7NewsBrisbane’ posts around the same amount as these accounts. The observed asymmetric patterns might hint at a centralised role for ‘7NewsBrisbane’ within the network, or it could be an artifact of an automated posting process (as [7] highlight) that takes longer than the 300 s window used for this study with ‘7NewsBrisbane’ coincidentally in the middle of list of accounts to post from. Further disentangling of these behaviours is an open question which is revealed but not resolved by the use of asymmetric networks. In this particular case where it is clear that the accounts form part of a news network for dissemination of news articles to slightly different geographic audiences, whether the coordination is automated or editorial is of not great interest: in cases where the intent and structure is not so clear it is an open question as to what can be inferred from situations like this, a topic we will return to in our discussion of our future plans.



**Fig. 6** Asymmetric Coordination Patterns in the 7News Subnetwork

### Mexico without plastics campaign

[7], Sect. 5.2.2] uncovered a coordinated “Mexico Without Plastics” campaign to prohibit single-use plastics in the country. The coordination occurred around a petition hosted by Greenpeace Mexico requesting national legal reforms as well as mentions of the official Mexican Senate account or its members. While the campaign is unrelated to the Covid lockdown measures, it is present in the dataset because of the term ‘libérate’ as in ‘free yourself from plastic’. The original study deployed a higher-order network of URL-mentions (where accounts are connected if they used the same URL and @mention in at least one tweet) and found a campaign network consisting of two clusters connected by three hub nodes, in a bow-tie shaped structure. The authors suggest that the three hubs in the centre of the network were the key leaders promoting the campaign. In our Multi-Behaviour network we find the same network, however we only discover the co-link dimensions of this campaign, given that we do not include mentions as a coordination relationship and do not employ higher-order behaviours. We surface four sub-networks associated with this campaign (usernames redacted for



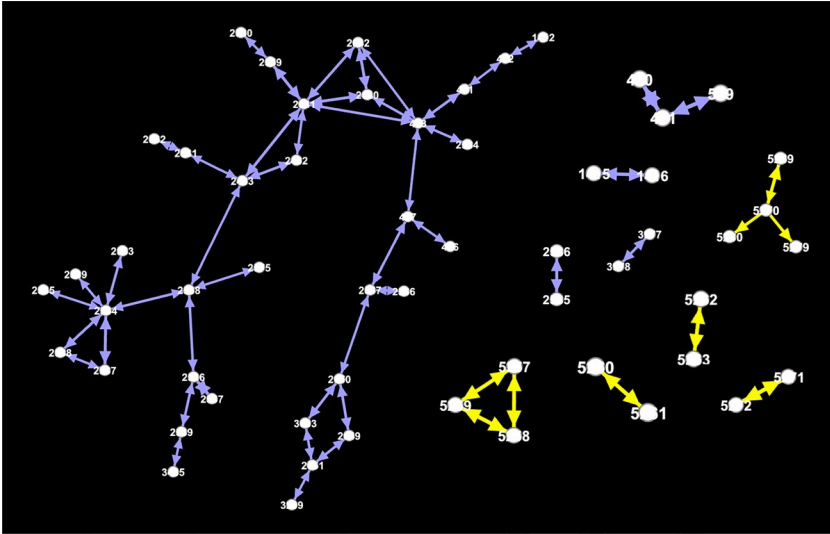


Fig. 7 The Open ICT for Indian Tax Inspectors campaign

privacy). It appears that the lack of @mentions relationships separates the key leaders into four distinct clusters rather than connecting them into a single connected cluster as found by [ [7], p. 11]. By keeping the different behaviours disjoint, i.e., not collapsed into a single edge denoting a higher-order relationship, we are able to observe variations in the scale and scope of coordination for this campaign. In this case, the network becomes disconnected if we do not consider @mentions and only focus on URL sharing. This demonstrates the usefulness of our approach, which provides analytic flexibility to capture a broader range of coordinated activity compared to other methods.

### Campaign to open ICT for Indian Tax Inspectors

The third campaign identified by [7] comprises Indian-based activists who appeal to various authorities and organisations to reopen the ICT (Inter Commissionerate Transfer), which would allow Inspectors in the CBIC (Central Board of Indirect Taxes and Customs) to have internal transfers. During the Covid pandemic, employees were separated from their families during lockdowns and experienced adversity. These activists formed a hashtag public to advocate for reopening of the ICT to address personal hardships caused by its closure, using hashtags such as #OpenICT-inCBICnow and #RestoreICTinCBIC to appeal to authorities and spamming @mentions to get attention. While [7] focus on higher-order hashtag-mention relationships through a single view, here we focus on any kinds of relationships detected through the Multi-Behaviour approach. As Fig. 7 shows, we discover a large co-retweet network consisting of 35 accounts as well as various clusters of coordinating pairs and triads. We identify two sub-campaigns where some accounts send duplicated messages (co-tweet; yellow edges) and others retweet the same content (co-retweet;

blue edges). The co-tweet sub-campaign introduce the #InspectorsWantICTinCBIC hashtag to voice their demands and focus more on narratives around mental health and stress compared to the co-retweet campaigners, although the co-tweeters also use the #OpenICTinCBIC hashtag in common with the co-retweeters. The co-tweeters also include links in their tweets whereas the co-retweets include @mentions. Our method thereby reveals different coordination strategies within the same campaign as well as disaggregating different behaviours in the same view, which provides more analytical depth compared to other methods. In this example, we not only capture the large co-retweet cluster but also multiple smaller clusters of co-retweet and co-tweet activity, which have interesting variations and suggest that our method captures subtle forms of coordination that qualitative analysis can explore further.

## Conclusion

In this paper we have introduced and evaluated an open-source software package and methodological framework for detecting and analysing coordinated behaviour on social media, namely the Coordination Network Toolkit. Our approach synthesises and improves upon the various methods for coordination analysis in the literature to offer a 'standardised' framework that increases the generalisability and usefulness of the methods. We suggest that it is important to establish community standards for interoperable evaluation of coordination networks: by releasing our toolkit as open source we hope it will be applicable and extensible by other researchers, and provide a common base for future work. This directly addresses a gap where there is currently no general purpose, open-source toolkit to construct coordination networks of various kinds using any kind of social media platform data. Our approach is thus multi-behavioural and multi-platform and is flexible so that the user can specify new types of relationships or behaviours as required. Compared to existing software packages, our toolkit we offer is highly optimised and can handle large datasets in the order of tens of millions of messages on standard desktop machines.

The methods we set out in this paper make several novel advances to the field that improve both analysis and visualisation of coordinated activity using social media data. We advocate for the use of weighted, directed multigraphs to more fully capture the structure and dynamics of coordinated activity across multiple behaviour types and in instances where there is asymmetry of participation for coordinating accounts. This helps to resolve the issue of false positives where many spammy accounts will perform the same behaviour as a target within the same time window, but the target never performs the same action as the spammers. The directed multigraph in large part helps analysts to differentiate and potentially filter out this activity where required, rather than erroneously collapsing it together in the undirected edges or simply removing such links as previous approaches have done. The use of edge labels and multiple edges to differentiate behaviour types allows analysts to visualise the complexity of coordination in a single dataset in a single view, and analyse coordination networks using standard graph algorithms that can handle directed and weighted edges. This provides some advantages compared to clustering

approaches for multi-level networks that require specialised algorithms and software or require multiple different views to focus on particular behaviours or higher-order action types.

To be sure, our approach is not intended to compete with other methods and frameworks but rather to complement and extend the field. We demonstrate this by reproducing the empirical case study of #ReopenAmerica as set out by [7] and showing the added benefits our method offers. The results of our analysis highlight two main takeaways. Firstly, our approach provides a ‘30000 foot’ of coordinated activity where multiple behaviour behaviours can be visualised and analysed in the same view simultaneously. In short, all the coordinated behaviour of various types is fully observed and traceable in the network. This differs from previous approaches that focus on single views in a clustered format. Secondly, while we discover the same campaigns in the #ReopenAmerica dataset as the original authors, the scale and scope of coordination that we detect offers a more flexible perspective. Our method identifies coordinated clusters that were not surfaced by previous work—these clusters are notable because of the complexity of their behavioural structure and dynamics, which is a product of our multidigraph approach. This shows the advantages of our approach in detecting new and diverse forms of coordination whilst also establishing its reliability in detecting similar to coordinated patterns in previous work.

## Limitations

While our approach provides a general method to detect multiple coordinated behaviour types simultaneously, we agree with [ [7], p. 9] that first-order types of behaviours can be limiting and lead to false-positive connections. For example, two users might share the same URLs repeatedly in a short time window, however they may do so due to opposing views (and thereby not coordinating at all). Future work and modifications to the Coordination Network Toolkit should implement higher-order behaviour types to capture more specific instances of coordinated activity, where accounts send posts that contain two or more behaviours simultaneously. For example, two posts may contain the same hashtags and also the same URLs, meaning that there are higher-order connections in the coordination activity. This is a key area for future work and will further enhance the accuracy of such methods at capturing genuine and not incidental or erroneous coordinated activity.

## Future work

Firstly, future work should implement higher-order behaviours into the Coordination Network Toolkit and evaluate how higher-order coordination analysis surfaces different and/or more accurate forms of coordinated activity. While the Multi-Behaviour approach we propose and evaluate in this paper captures all forms of observed coordinated behaviour, in doing so it may also capture erroneous forms of coordination that are coincidental rather than highly synchronised. Similarly, while our paper provides a baseline empirical comparison of results

compared to previous work, no studies have yet conducted a thorough comparative study of different coordination network analysis techniques to assess accuracy and reliability. We recommend that future studies undertake a more in-depth comparative analysis of different methods to better understand the limitations and advantages of different approaches to studying coordinated online behaviour.

Secondly, there is currently no statistical framework purposely designed to analyse coordination networks. For instance, we have little statistical understanding of what constitutes strong versus weak coordination and how the empirical distributions of coordinated patterns vary by context and platform. The research by [16] provides a useful departure point for future work. In this paper, [16] proposes a Poisson mixture model based on time intervals between shares of URLs on Facebook, to detect and verify the presence of coordinated activities between pairs of accounts in a dataset. For each pair of accounts, a set of inter-arrival times between URL shares and their associated log-odds ratios is compiled. A t-test is then applied to evaluate the hypothesis that the average log-odds ratio of coordinated link sharing surpasses zero. In other words, if the sharing of URLs appears more coordinated than random. Only edges between entities that fulfill a significance threshold of  $p < 0.05$  are retained. Importantly, the study adjusts for the risk of type I errors from multiple comparisons by implementing the Holm-Bonferroni procedure. This approach seems promising for further work to develop statistical frameworks to deepen the theory and practice of coordination network analysis.

Finally, we agree with [7] that future work should examine intent discovery, whereby conceptual and technical approaches are developed to better understand why accounts appear to coordinate on particular issues or topics. Related to this is the conceptualisation of coordination, which appears to be in its infancy. Various studies highlight the slipperiness of this term, whereby incidental or accidental co-occurring timing of online behaviours is often mixed in with intentional coordination. We therefore suggest that future work develop a stronger conceptual basis for coordination network analysis, which is able to meaningfully and accurately distinguish various forms of online participation and synchronisation.

**Acknowledgements** The authors acknowledge support from the Queensland University of Technology (QUT) Digital Observatory, the Australian Research Council (ARC) Centre of Excellence in Automated Decision-Making + Society (ADM+S), and the Digital Media Research Centre (DMRC) at QUT. We also acknowledge the contribution of Guangnan Zhu, who conducted early prototyping and testing of the coordination network analysis methods.

**Author Contributions** TG designed the research study, with input from all authors. TG and SH wrote the paper with input from BH. TG conducted the empirical analysis. SH and BH programmed the original open-source toolkit, supervised by TG. GZ conducted early prototyping and testing of the toolkit.

**Funding** Open Access funding enabled and organized by CAUL and its Member Institutions. This research is funded by an ARC Discovery Early Career Research Award (DECRA) grant: DE220101435.

**Data availability** The Twitter dataset is not available for public release because it violates the Terms of Service and contains identifying information. An anonymised dataset is available upon request and tweet IDs are also available upon request.

**Code availability** All code to reproduce the analysis in this paper is available upon request from the leading author. The open-source software in this study is available at: <https://github.com/QUT-Digital-Observatory/coordination-network-toolkit>.

## Declarations

**Conflict of interest** On behalf of all authors, the corresponding author states that there is no conflict of interest.

**Ethics approval** The study is approved for ethics by the QUT Human Research Ethics committee as a Human Ethics Low Risk Application (approval: LR 2023-5307-15435).

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

1. Badawy, A., Ferrara, E., & Lerman, K. (2018). Analyzing the digital traces of political manipulation: The 2016 Russian interference Twitter campaign. In *2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pp. 258–265. IEEE.
2. Stewart, L. G., Arif, A., & Starbird, K. (2018). Examining trolls and polarization with a retweet network. In *Proceedings of WSDM Workshop on Misinformation and Misbehavior Mining on the Web (MIS2)*, p. 6. ACM, New York, NY, USA. [https://doi.org/10.475/123\\_4](https://doi.org/10.475/123_4)
3. Graham, T., & Observatory, Q. D. (2020). Coordination Network Toolkit. Queensland University of Technology. (Software). [https://doi.org/10.25912/RDF\\_1632782596538](https://doi.org/10.25912/RDF_1632782596538)
4. Weber, D., & Neumann, F. (2021). Amplifying influence through coordinated behaviour in social networks. *Soc. Netw. Anal. Min.*, *11*, 111. <https://doi.org/10.1007/s13278-021-00815-2>
5. Giglietto, F., Righetti, N., Rossi, L., & Marino, G. (2021). Coornet: An integrated approach to surface problematic content, malicious actors, and coordinated networks. *AoIR Selected Papers of Internet Research*.
6. Pacheco, D., Flammini, A., & Menczer, F. (2020). Unveiling coordinated groups behind white helmets disinformation. In *Companion Proceedings of the Web Conference 2020*, pp. 611–616.
7. Magelinski, T., Ng, L., & Carley, K. (2022). A synchronized action framework for detection of coordination on social media. *Journal of Online Trust and Safety* *1*(2).
8. Keller, F. B., Schoch, D., Stier, S., & Yang, J. (2020). Political astroturfing on twitter: How to coordinate a disinformation campaign. *Political Communication*, *37*(2), 256–280.
9. Schafer, F., Evert, S., & Heinrich, P. (2017). Japan's 2014 general election: Political bots, right-wing internet activism, and prime minister Shinzō Abe's hidden nationalist agenda. *Big Data*, *5*(4), 294–309.
10. Schoch, D., Keller, F. B., Stier, S., & Yang, J. (2022). Coordination patterns reveal online political astroturfing across the world. *Scientific Reports*, *12*(1), 4572.
11. Giglietto, F., Righetti, N., Rossi, L., & Marino, G. (2020). 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.
12. Massanari, A. (2017). # gamergate and the fapping: How reddit's algorithm, governance, and culture support toxic technocultures. *New Media & Society*, *19*(3), 329–346.
13. Cruickshank, I. J. (2020). Multi-view clustering of social-based data. PhD thesis, Carnegie Mellon University.
14. Bastian, M., Heymann, S., & Jacomy, M. (2009). Gephi: an open source software for exploring and manipulating networks. In *Proceedings of the International AAAI Conference on Web and Social Media*, vol. 3, pp. 361–362.

15. Jacomy, M., Venturini, T., Heymann, S., & Bastian, M. (2014). Forceatlas2, a continuous graph layout algorithm for handy network visualization designed for the gephi software. *PloS one*, 9(6), 98679.
16. Broniatowski, D. A. (2021). Towards statistical foundations for detecting coordinated inauthentic behavior on facebook. Techreport Preprint, Institute for Data, Democracy and Politics-The George Washington University. Last accessed on, 05–14.

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

## Authors and Affiliations

Timothy Graham<sup>1</sup>  · Sam Hames<sup>2</sup> · Elizabeth Alpert<sup>3</sup>

✉ Timothy Graham

Timothy.Graham@qut.edu.au

<https://scholar.google.com.au/citations?user=9GG-wWEAAAAJ>

Sam Hames

s.hames@uq.edu.au

<https://scholar.google.com.au/citations?user=rwXGu00AAAAJ>

Elizabeth Alpert

elizabeth.alpert@qut.edu.au

<sup>1</sup> School of Communication, Queensland University of Technology, Victoria Park Road, Kelvin Grove 4059, QLD, Australia

<sup>2</sup> School of Languages and Cultures, The University of Queensland, Brisbane 4072, QLD, Australia

<sup>3</sup> Digital Observatory, Queensland University of Technology, George Street, Brisbane 4000, QLD, Australia