September 2023

A CONSIDERATION OF COORDINATED BEHAVIOUR IN THE #AUKUS CONVERSATION ON X (FORMERLY TWITTER)

Timothy Graham

Queensland University of Technology, Digital Media Research Centre

Darren L. Linvill

Clemson University, Watt Family Innovation Center Media Forensics Hub

Patrick Warren

Clemson University, Watt Family Innovation Center Media Forensics Hub

Katherine M. FitzGerald

Queensland University of Technology, Digital Media Research Centre

ABOUT CLEMSON

The Media Forensics Hub is an interdisciplinary team of researchers working to study and combat online deception with the goal of building society's resilience to the dangers it poses. The Media Forensics Hub was launched in May 2020 with support from the South Carolina Research Authority as a project of Clemson's Watt Family Innovation Center. In 2022, the Hub received major support from the John S. and James L. Knight Foundation with the goal of expanding their impact. The Hub has partnered with a range of organizations with the goal of identifying and mitigating the impact of disinformation. These include the Commission on Presidential Debates, CNN, and ProPublica to name a few.

ABOUT QUT

The Queensland University of Technology's Digital Media Research Centre (DMRC) is one of the world's leading institutes for digital humanities and social sciences research. Our research pioneers innovative digital methods for social and cultural research, in order to help governments, communities and businesses understand and adapt to the changing digital media landscape. QUT has been consistently ranked #1 in Australia and was ranked #19 overall worldwide for Communication and Media Studies in the recent QS World University Rankings, 2023. The research underpinning the technology in this report was invented by Associate Professor Timothy Graham and supported by funding from the Australian Research Council through the DECRA scheme.

03

Executive Summary

This case study is a collaboration between researchers at Clemson University, United States and Queensland University of Technology, Australia who collectively have over two decades of complementary expertise in the detection of inauthentic actors in contested digital spaces.

In this case study, we demonstrate a unique combined approach to detecting coordinated activity on social media to a degree of accuracy and sensitivity that is not yet seen anywhere in the market. The approach we bring to bear - combining network science with statistical anomaly detection around the objects of influence - not only detects the presence of networks of coordinated inauthentic accounts, but also reveals many details about their strategies and the objects of influencing that they are targeting. This detailed intelligence can have enormous strategic value in understanding where adversaries are allocating their efforts, for understanding the strengths and weaknesses of their approaches to influence, and for planning countermeasures.

This project analyzed 256,655 tweets containing the hashtag #AUKUS or #NoAUKUS collected in March 2023. Our process is flexible enough to be used across platforms for cross-platform coordinated inauthentic campaign detection, but as it was developed on Twitter (now known as X), we present an application in that context.

We identified seven substantial coordinated influence operations active on this, quite narrow, topic. They varied enormously in size, scope, and sophistication, as well as the degree to which they were deceptive about their purpose. At one extreme, we found a small network of a dozen related accounts with a clear narrow mission, which mostly disclosed their origin and purpose. At the other extreme, we discovered a very large network of clearly coordinated inauthentic accounts that are pushing narratives in support of India (especially the Indian military) and against Pakistan and China. There is also a very large network of highly coordinated relatively authentic accounts, pushing narratives and promotion properties related to exiled Chinese billionaire Miles Guo.

Importantly, the approach empowers analysts to work efficiently and with powerful breadth of context to pull from. The approach enables efficient detection and contextualisation of online influence campaigns and inauthentic behavior. For the first time, reliable and detailed intelligence can be provided in a timely manner (hours or days rather than weeks and months) that leaves nowhere for actors to hide, reveals tactics and distinguishes genuine from malicious activity.

04

Glossary

- **Baby:** Tweet or post produced by an account with fewer than 100 tweets.
- **Co-linker:** Tweet or post produced by an account with a co-link degree of more than 20. The degree is the total sum of edges to and from a node. This means that the account has linked the same URL as other accounts more than 20 times.
- **Egg:** Tweet produced by an account with no description in their profile.
- **First:** A tweet or post that is produced in the first second of the minute. For example, a post at 13:00:01. This is an indicator of automation as all scheduled or automated posts will be sent within the first applicable second.
- **Flood:** A tweet or post that has more than 10 verbatim copies by more than two distinct accounts, none of which are verified. The rationale behind this is that non-verified and potentially inauthentic accounts are "flooding" the conversation with repetitive messages in an attempt to hijack the conversation and narrative.
- **Nodes and Edges:** For the purpose of this report, nodes in graphs represent social media accounts or users. Edges visualize the links between these users. The 'weight' of edges can be heavier depending on the strength of the relationship between two actors. This is represented by thicker lines in network visualisations.

05

Introduction

The AUKUS Partnership

Announced in September 2021, AUKUS is a trilateral security partnership between the United States, Australia, and the United Kingdom (Kahn, 2023). The partnership aims to enhance security and stability in the Indo-Pacific region by sharing advanced technology, including nuclear-powered submarines (U.S. Embassy in Canberra, 2023). The AUKUS submarine debate refers to the controversy that arose over Australia's decision to cancel a contract with France for the supply of 12 French diesel-powered submarines, and instead opt for the nuclear-powered submarines provided by the US and UK (Murphy & Hurst, 2021). The French government expressed its disappointment and anger over the decision, which it saw as a breach of trust that jeopardized strategic partnerships in the future (Hurst, 2021).

In March 2023, the government announced that Australia will purchase a minimum of three nuclear powered Virginia-class submarines from the US to arrive in the early 2030s (Ali et al., 2023; Prime Minister of Australia, 2023). The decision by Australia to choose nuclear-powered submarines has been criticized by some experts who believe that such submarines are not necessary for Australia's defense needs and could escalate tensions in the region, particularly with China, who have been vocally critical of the deal in foreign ministry press briefings (Pengelly & Hawkins, 2023).

The AUKUS submarine debate has raised questions about the role of nuclear weapons in regional security and the importance of maintaining good relations with allies. It has also highlighted the complex geopolitical challenges facing countries in the Indo-Pacific region, which are grappling with rising tensions and competing interests. The AUKUS submarine deal is a sensitive foreign affairs topic that has global ramifications given the increased militarisation of Australia and its position as a critical US ally in the Indo-Pacific. The Twitter debate around the submarines is an ideal candidate for information warfare and foreign threats in a way that targets public opinion and, in doing so, shapes geopolitics.

Other countries in the Asia Pacific region, such as China, have a vested interest in this contested military and digital space. China has previously expressed its disapproval of the AUKUS.

06

In a tweet on 14th of March 2023, Beijing's mission to the United Nations stated: "The nuclear submarine cooperation plan released today by #AUKUS is a blatant act that constitutes serious nuclear proliferation risks, undermines international non-proliferation systems, fuels arms races, and hurts peace and stability in the region". Given these hard-line statements against the security partnership, we aim to investigate potential digital propaganda or foreign influence in the debate.

Our Synthetic Approach

This report is a collaboration between Queensland University of Technology (QUT) and Clemson University (CU). Both partners in the collaboration analyzed the same dataset, applied their own methods for detection of coordinated inauthentic behavior, and we synthesized these results in this report. Greater integration of these methods will allow for even more complete detection of coordination and offer analysts greater context and depth of understanding than is available on current social media listening platforms.

The two broad streams that make up our synthetic approach to the detection of coordination are quite different but entirely complementary. The first stream (from QUT), identifies coordination by looking at patterns of timing in online behaviors. Specifically, two social media accounts are linked if they perform the same actions repeatedly at very close time intervals. For example, accounts are linked if they share the same URL repeatedly within a 60 second time window, or if they both retweeted the same tweet repeatedly in the same window. The strength of the coordination between accounts increases as pairs of accounts engage in more and more of this temporally-aligned activity. Clusters are then formed from groups of pairs of coordinating accounts, thereby revealing both authentic and inauthentic collective action. Typically, unusual coordination patterns (e.g. extremely strong and densely clustered connections) are giveaways for coordinated inauthentic activity.

The second stream (from Clemson), identifies coordination by looking for unusual patterns of accounts' characteristics across sets of messages that share similar influence targeting. Specifically, a set of messages becomes suspicious if the set of accounts that are contributing to that set are an outlier along some set of account characteristics. For example, if an unusual share of accounts sharing some hashtag have been created in the last 7 days, we become suspicious of both those accounts and of that hashtag. To the extent that commonality is shared across many hashtags we become increasingly suspicious of them all.

This report is a collaboration between Queensland University of Technology (QUT) and Clemson University (CU). Both partners in the collaboration analyzed the same dataset, applied their own methods for detection of coordinated inauthentic behavior, and we synthesized these results in this report. Greater integration of these methods will allow for even more complete detection of coordination and offer analysts greater context and depth of understanding than is available on current social media listening platforms.

The two broad streams that make up our synthetic approach to the detection of coordination are quite different but entirely complementary. The first stream (from QUT), identifies coordination by looking at patterns of timing in online behaviors. Specifically, two social media accounts are linked if they perform the same actions repeatedly at very close time intervals. For example, accounts are linked if they share the same URL repeatedly within a 60 second time window, or if they both retweeted the same tweet repeatedly in the same window. The strength of the coordination between accounts increases as pairs of accounts engage in more and more of this temporally-aligned activity. Clusters are then formed from groups of pairs of coordinating accounts, thereby revealing both authentic and inauthentic collective action. Typically, unusual coordination patterns (e.g. extremely strong and densely clustered connections) are giveaways for coordinated inauthentic activity.

The second stream (from Clemson), identifies coordination by looking for unusual patterns of accounts' characteristics across sets of messages that share similar influence targeting. Specifically, a set of messages becomes suspicious if the set of accounts that are contributing to that set are an outlier along some set of account characteristics. For example, if an unusual share of accounts sharing some hashtag have been created in the last 7 days, we become suspicious of both those accounts and of that hashtag. To the extent that commonality is shared across many hashtags we become increasingly suspicious of them all.

The AUKUS Dataset

We collected 256,655 tweets sent by 75,079 unique accounts between 1st March 2023 and 27th March 2023 containing the key terms 'AUKUS' or 'nuclear powered submarine' or 'nuclear submarine', or the hashtags #AUKUS or #NOAUKUS that represented the polarized stances in support of and opposition to the trilateral agreement. We refer to this as the "AUKUS debate dataset". Our goals are to identify sets of accounts that are contributing to this topic in a coordinated fashion, to delineate what elements of the topic they are targeting, to infer as much as we can about their tactics and strategies, and to infer who is behind them.

08

Findings

Below we outline findings from the AUKUS debate dataset. Six interesting networks were identified in this dataset. Some of these networks were relatively innocent, authentic actors. Others are likely state affiliated influence operations. The tools we have developed offer context for understanding differences between networks. We present findings in order from the network that is subjectively the most authentic (9News) to that which is subjectively least authentic (possible Indian state affiliated troll network).

NOTE: Most of the networks we describe below were, in fact, identified using more than one method of analysis. While this fact itself is extremely valuable for triangulation purposes during analysis, for simplicity's sake in some cases we will be discussing only one method of identification per network.

9News Network: Eight Coordinated Accounts Identified

In examining popular hashtags used in the AUKUS debate dataset, we find several that have suspicious mixes of flagged accounts. On two separate days in the data, #9News was used at a high rate by accounts with the “first” flag, meaning accounts which were likely to tweet in the first second of the minute (a sign of automation).

first

hashtags

	convo	first	tweet_count	user_count	short_day
1	#FOI👇	0.090909	44	44	03/12/2023
3	#FreeJulianAssangeNOW	0.057692	52	4	03/15/2023
5	#UnitedStates	0.066667	60	45	03/14/2023
6	#9News	0.058824	51	39	03/13/2023
8	#9News	0.209302	43	29	03/15/2023
8	#Biden	0.064935	77	72	03/14/2023
8	#news	0.063830	47	31	03/14/2023

Figure 1: A list of hashtags associated with the ‘first’ bot indicator flag, with #9News highlighted.

09

Analysis of these accounts' behavior also suggests they are coordinated. These accounts coordinate the sharing of news articles and are thus strongly co-linking (red edges in Figure X) with each other, however we also observe co-retweet (blue) edges directed out from these accounts to various peripheral Twitter accounts that are highly active in news sharing on Australian Twitter. In this case, the 9News accounts are behaving in two types of coordinated behavior.

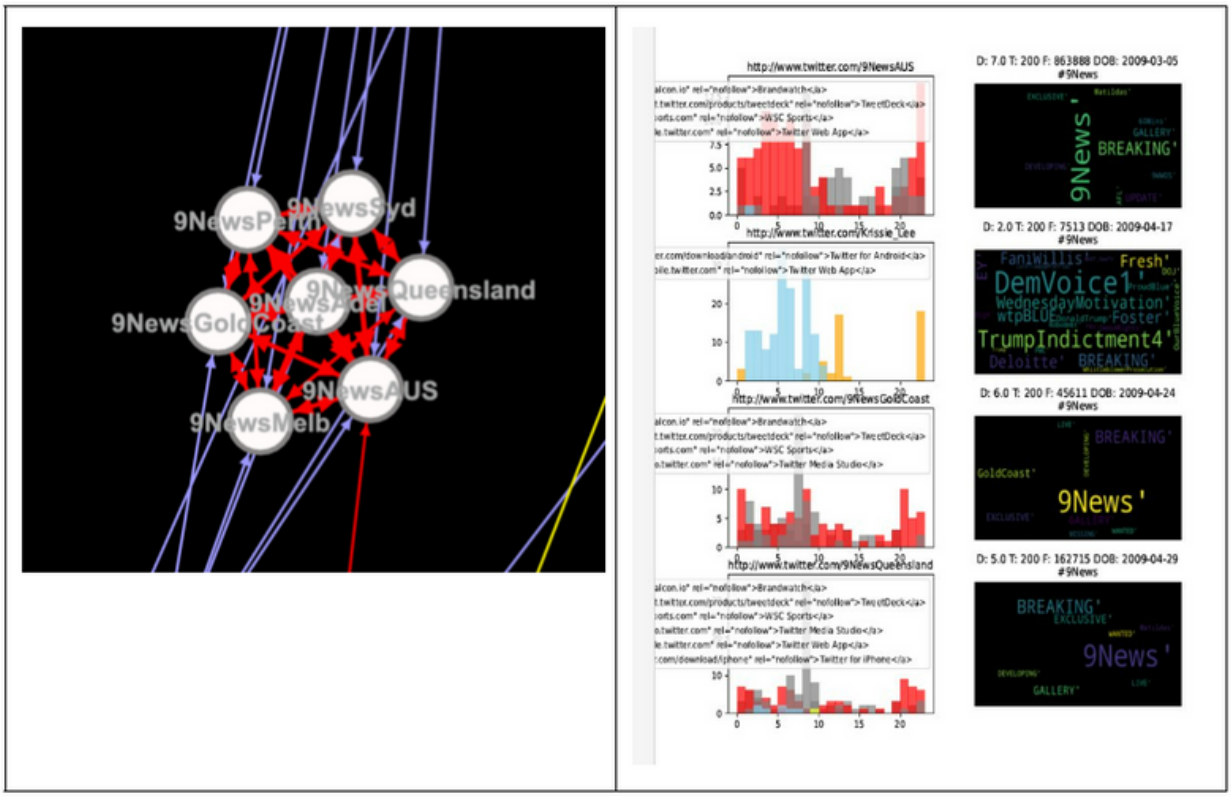

Figure 2: Histograms, networks, and word clouds for four 9News accounts. The color of the histograms indicates the posting client.

Finally, we have developed a tool for visualizing account behavior in a manner that not only gives quick insight into the conduct of the account but allows an analyst to make more refined judgements regarding coordination. Below are a set histograms for four accounts with a “first” flag which used #9News on a given day. The histogram shows the past 200 tweets from each account and the hour of the day in which each tweet occurred. The color further indicates the client (e.g., Twitter for Web App, Twitter for Android, Twitter for iPhone, etc.) employed. Along with each hashtag is a word cloud of the prominent hashtags employed by the account, giving some idea of what the account said in those tweets. Three out of four of the accounts shown in these histograms belong to 9News and are clearly operated in a manner that is very similar to one another: they tweet in an automated manner throughout the day, employ the same clients, and use very similar hashtags to one another. This visual analysis allows us to quickly remove any accounts from suspicion of coordination if they operate in a manner inconsistent with the other accounts.

The 9News accounts are clearly an authentic coordinated network. They are operated by a legitimate news network and represent their organization honestly. Our process allows us to see and understand this context.

Figure 3: Screenshots of two of the 9News outlet profile bios, demonstrating they are an authentic media outlet.

11

Kidman Network: 12 Coordinated Accounts Identified

We analyzed hashtags that are outliers as a results of the large number of 'baby' accounts engaging with them; 'baby' accounts being those that have produced fewer than 100 tweets. In doing so, we find two hashtags on two days that are outliers due to the same set of babies: #AUKUS and #defence.

baby

hashtags

	convo	baby	tweet_count	user_count	short_day
1	#AUKUSAlbo	0.071429	42	38	03/14/2023
1	#OTD	0.056604	53	53	03/09/2023
4	#AUKUS	0.065657	198	164	03/05/2023
5	#UnitedStates	0.216667	60	45	03/14/2023
6	#UnitedKingdom	0.060000	50	42	03/14/2023
7	#australia	0.063830	47	38	03/14/2023
7	#defence	0.239130	46	37	03/12/2023
13	#Russia	0.067797	59	54	03/15/2023
14	#Ukraine	0.071429	56	52	03/15/2023
15	#Russia	0.056604	53	50	03/17/2023
18	#US	0.071429	42	37	03/22/2023

Figure 4: A list of hashtags that are associated with the 'baby' bot indicator flag, with #AUKUS highlighted.

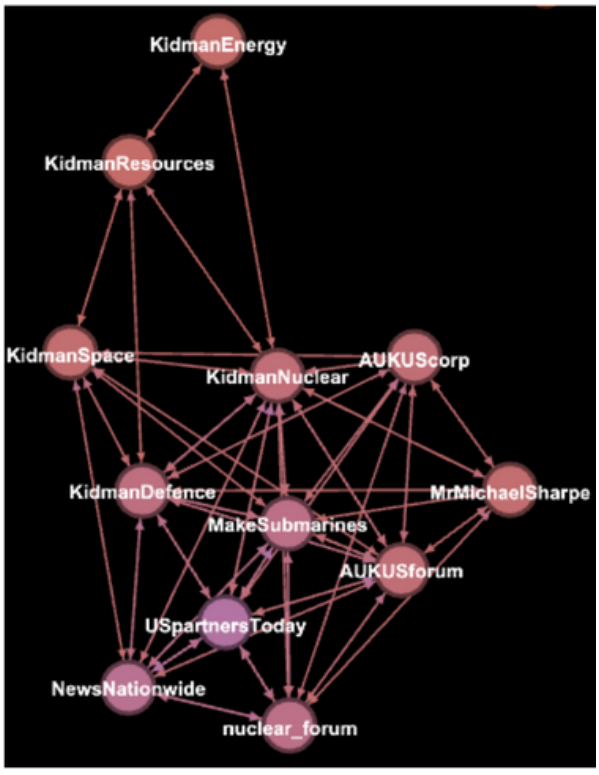

Figure 5: Network of 'Kidman accounts' associated with Michael Sharpe, leader of the AUKUS Forum.

There are twelve of these accounts, all of which operate in the same manner. Analysis of their behavior shows they tweet at consistent times, on similar topics, and with the same client (Twitter for iPhone). These accounts have few followers and limited engagement. In their profile information, most of these accounts self-identify as being affiliated with Michael Sharpe, an Australian business person and founder of the AUKUS Forum, an organization specifically dedicated to supporting the AUKUS partnership.

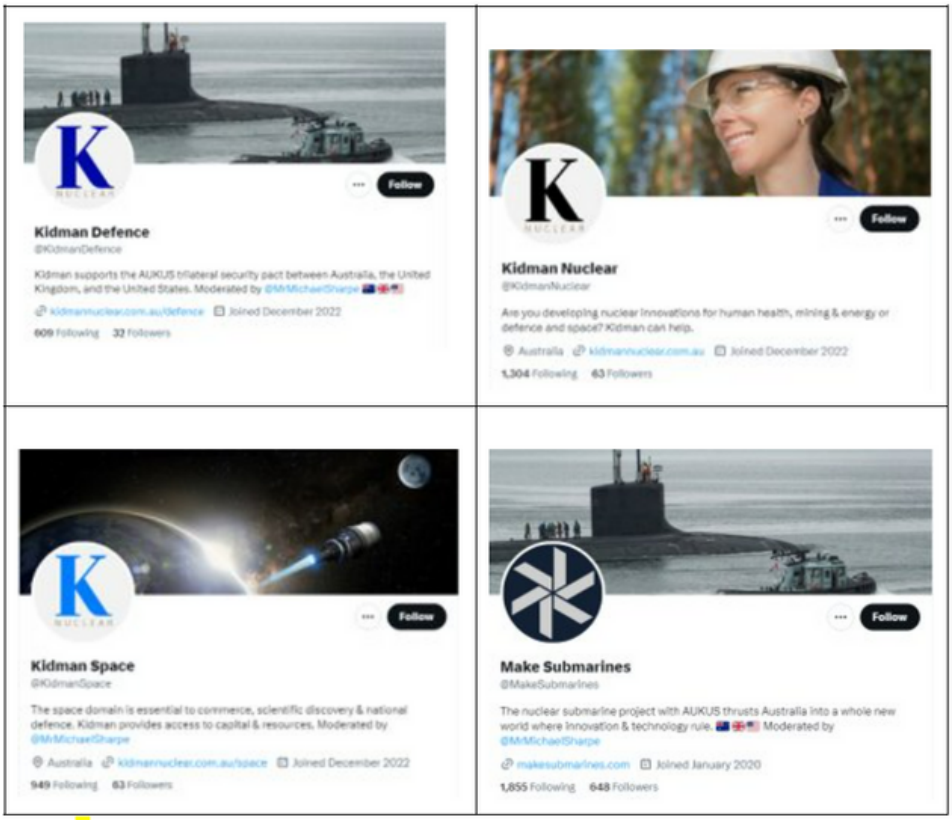

Figure 6: Collection of accounts that are affiliated with Michael Sharpe, leader of the AUKUS Forum.

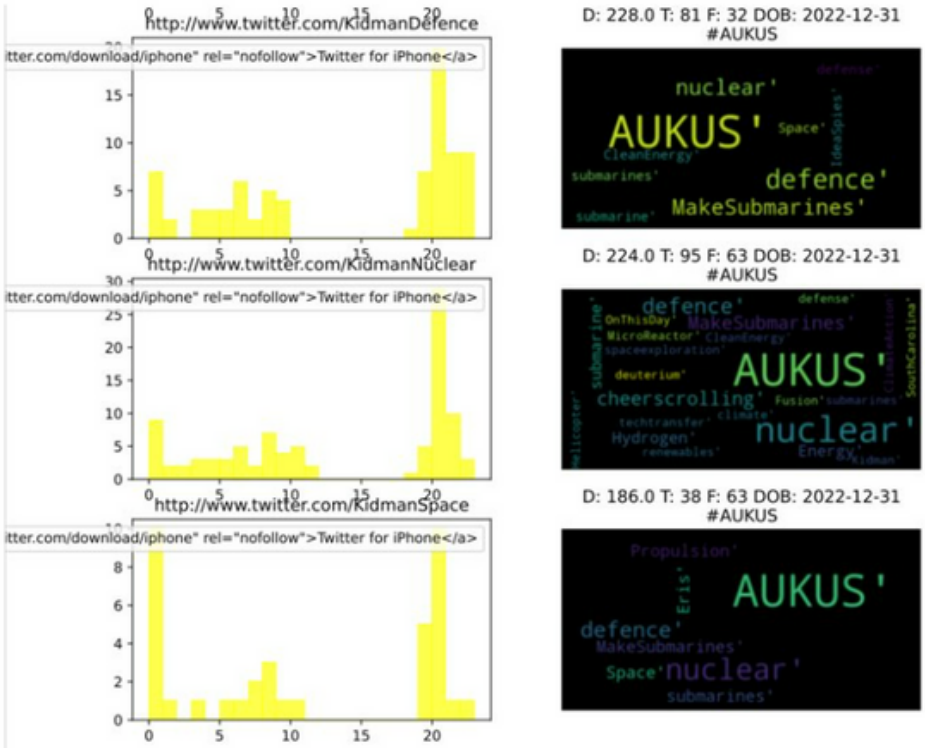

Figure 7: Histogram and word cloud of three of the recent posts of accounts associated with Michael Sharpe: KidmanDefence, KidmanNuclear, and KidmanSpace.

One account in this group, however, appears as an exception. @NewsNationwide is clearly operated by the same actor as it follows identical patterns. It, however, claims no connection to Michael Sharpe and presents itself as an unaffiliated news aggregator. Interestingly, this account is somewhat older and with more followers than the other accounts in the network.

Figure 8: Histogram and word cloud of three of the recent posts of the account @NewsNationwide.

The coordination analysis approach also discovers the Kidman network, as shown in the network diagram above in Figure X. While the previous analysis found seven accounts, the coordination approach adds to this by identifying an additional four accounts: @AUKUScorp, @AUKUSforum, @nuclear_forum, @USpartnersToday, along with showing the coordination with Michael Sharpe's account (@MrMichaelSharpe). These accounts frame the AUKUS deal as the most substantial upgrade to Australia's military capability since World War II, presenting a significant industrial challenge. They coordinate to advertise and organize events that aim to connect industry professionals, researchers, and national security personnel. These events bolster the AUKUS initiative and emphasize the submarine manufacturing aspect of the security agreement, along with the importance of technological innovation and partnership against the backdrop of global power dynamics. The 9News accounts are clearly an authentic coordinated network. They are operated by a legitimate news network and represent their organization honestly. Our process allows us to see and understand this context.

Figure 9: Profiles of four additional accounts within the 'Kidman Network'.

15

Miles Guo Network: 204 Coordinated Accounts Identified

Turning from examining hashtags as the object of influence to domains, we find several domains pushed by unusually high numbers of “eggs” (accounts using the default profile picture).

domains

	convo	egg	tweet_count	user_count	short_day
0	www.vtv.gob.ve	0.539514	658	658	03/14/2023
0	www.himalayaustralia.com.au	0.505882	85	85	03/06/2023
0	www.zerohedge.com	0.502538	197	196	03/10/2023
1	www.vtv.gob.ve	0.562604	599	554	03/15/2023
1	www.himalayaustralia.com.au	0.490909	110	110	03/07/2023
1	www.rfi.fr	0.502326	215	215	03/14/2023
1	www.voachinese.com	0.557692	52	52	03/10/2023
2	www.firstpost.com	0.676923	65	65	03/15/2023
4	thediplomat.com	0.666667	54	54	03/15/2023
6	www.himalayaustralia.com.au	0.494949	99	99	03/12/2023
7	www.himalayaustralia.com.au	0.495726	117	117	03/13/2023
8	www.youtube.com	0.467742	62	62	03/10/2023
15	www.msn.com	0.542553	188	187	03/17/2023

Figure 10: A list of domains that are being shared by a disproportionate number of ‘eggs’, accounts using the default profile picture.

Several distinct networks are promoting different domains. Looking closely, however, we find that the same, large coordinated set of accounts is promoting stories appearing on himalayaustralia.com, zerohedge.com, rfi.fr, and voachinese.com. Looking closely at these accounts, we observe some variety in how they behave. There are subsets within the group that have high degrees of similarity, but there is, overall, some heterogeneity. All of the accounts focus on the same topics of conversation, however, and still demonstrate some coordination. These accounts are part of the ‘Whistleblower Movement’ or ‘Whistleblower Revolution’ (爆料革命), a campaign funded by billionaire Guo Wengui (Miles Guo), a Chinese businessman in exile in the US. See the appendix for further information regarding Miles Guo and the Whistleblower Movement. The network is made up of real people, likely with genuine beliefs. They are, however, highly coordinated and it seems probable, based on our analysis, that single users operate multiple accounts.

One subset of accounts discovered in this cluster goes by the name '澳喜特战旅' (Aussie Special Brigade). This subnetwork consists of 40 nodes and 78 edges, however, there are many other smaller subnetworks of fake accounts that are part of this campaign spread throughout the entire bot-like network (e.g., dyads and triads of coordinating accounts).

Figure 11: Profile, histogram, and word cloud of three accounts associated with the Miles Guo cluster.

Figure 12: Example tweets shared from accounts in the Miles Guo network.

The coordination analysis approach also discovers the Miles Guo network, which turns up as a densely connected cluster of 204 Twitter accounts as shown in the network diagram below. We observe the complementarity of the coordination analysis approach as it detects an additional 33 Twitter accounts that are part of this network. The accounts are connected because they repeatedly retweet the same content within 60 seconds of each other. These 204 accounts are connected by 5074 edges or links in the network.

Venezuelan Astroturfers: 559 Coordinated Accounts Identified

Also among the domains promoted by "egg" accounts is a Venezuelan domain: vtv.gob.ve. This domain is being promoted by a large number of accounts. These accounts have several behavioral patterns that are suspicious. Most notably, they post at extremely high rates for short periods of time and engage exclusively on topics supportive of the Venezuelan government. Many accounts almost exclusively retweet Venezuelan state media and government accounts. These accounts are likely part of a network maintained by the Venezuelan government. The Maduro government uses money and prizes to encourage Venezuelans to post on behalf of the regime. These accounts appeared in our dataset because they reposted a story from Venezuelan state media about AUKUS.

Figure 13: An example of content promoted by the Venezuelan astroturfing network

18

Marketing Accounts: 28 Coordinated Accounts Identified

Social media accounts operated for marketing or public relations purposes are commonplace, and this dataset was no different. One set of "egg" accounts was found to be promoting links to the youtube.com domain. Upon examination, many of these accounts engaged in very similar behavior to one another. This network all employed Twitter for Android, tweeted throughout the day (with no clear human sleep cycle), were anonymous, seemingly had no organic followers, and engaged almost exclusively in retweeting and liking content. Most importantly, however, this set of accounts also all focused their engagement around @DokterTifa, an account belonging to Tifauzia Tyassuma, an Indonesian medical doctor, author, and social media influencer.

@DokterTifa shared posts containing links to YouTube.com and which were critical of AUKUS. A coordinated network working to promote the account retweeted these posts as they would any other post from @DokterTifa.

Figure 14: The Twitter profile of Tifauzia Tyassuma, aka @Dokter Tifa, an Indonesian doctor and social media influencer.

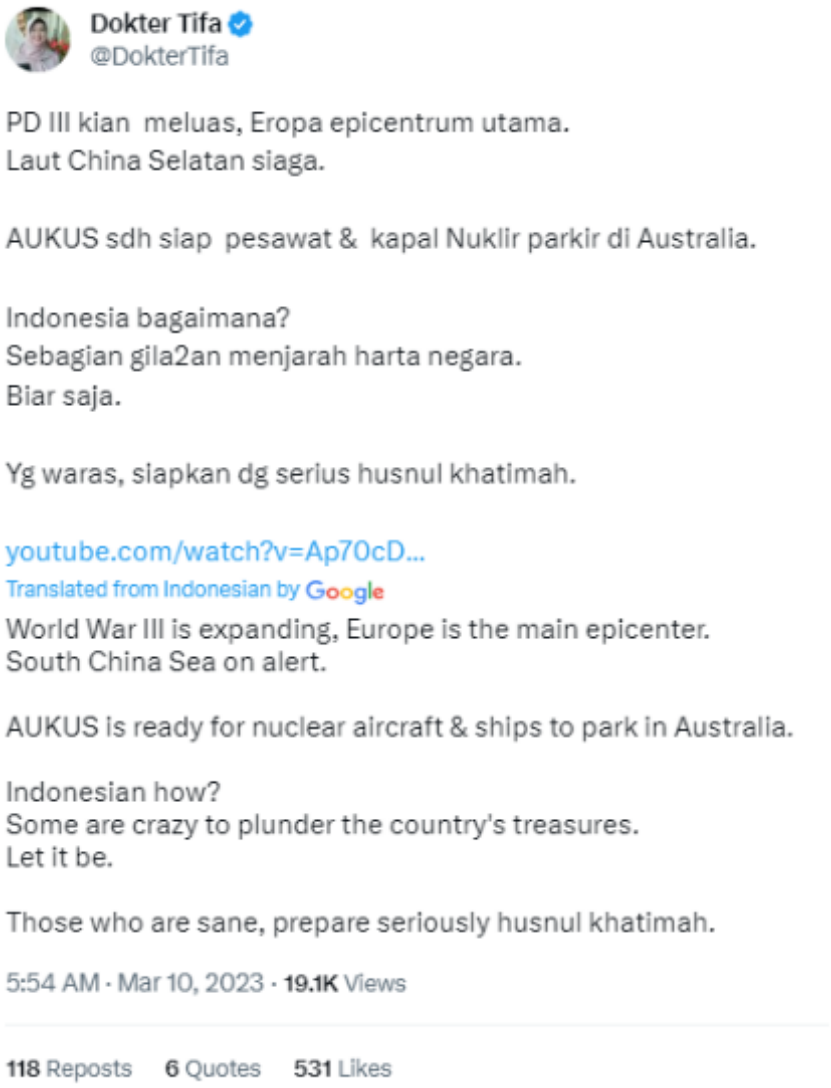

Figure 15: An example tweet from @DokterTifa's account.

Pro-Indian Network #1: 312 Coordinated Accounts Identified, and linked to a larger network of over 600

By examining domains that were mentioned by a large share of accounts with the "flood" flag (tweets that are exactly duplicative) one last large coordinated set of accounts, this one quite large. Accounts promoting links to asiatimes.com, indiatimes.com msn.com, wsj.com, the diplomat.com, and firstpost.com all engage with strikingly similar behavior. The timing of their tweets is systematic, they are all anonymous, many have identical creation dates, and their content is consistent across accounts. Interestingly, this network seems to be operated in several distinct groups, some employing Twitter for Android, others using Twitter Web App, and a number that use both (but in the exact same pattern).

20

These accounts all have a clear pro-Indian government agenda. They are highly critical of both Pakistan and China (to an extreme degree) and also generally pro-Western. They share links to news and opinion as well as political cartoons. The level of engagement the accounts receive is not what one might expect from a real human user, but they all receive at least some views. These accounts shared links to articles supporting AUKUS, particularly its indirect benefits to India.

Figure 16: Profile, histogram, and word cloud of three accounts associated with the pro-Indian cluster.

Figure 17: Profile, histogram, and word cloud of three accounts associated with the pro-Indian cluster.

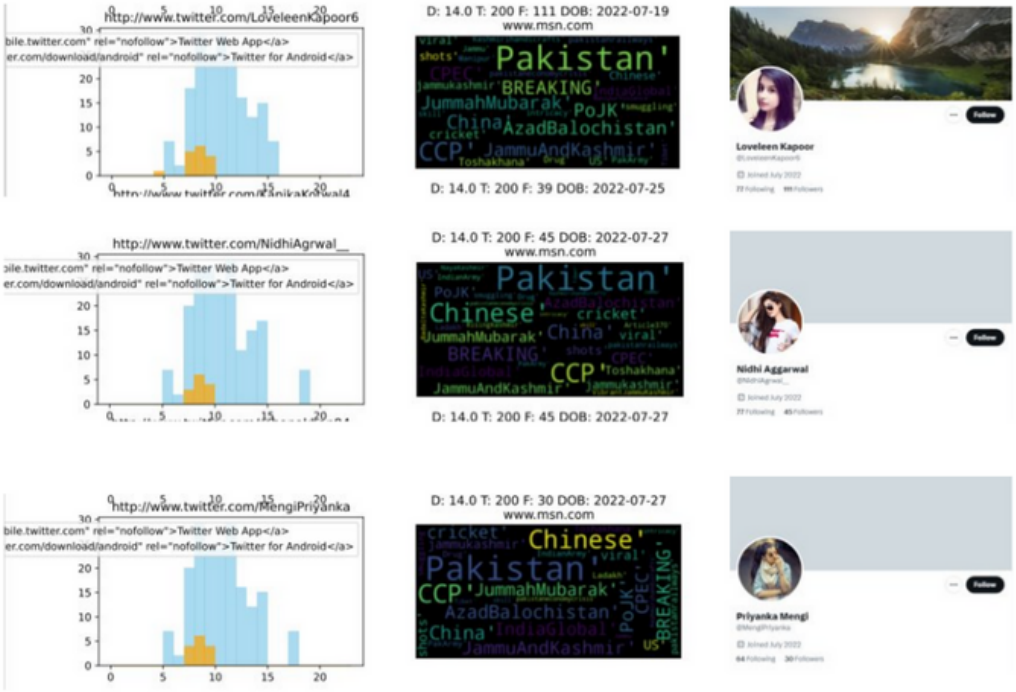

Figure 18: Profile, histogram, and word cloud of three accounts associated with the pro-Indian cluster.

There is substantial evidence that these accounts are centrally directed and organized but operated by a collection of decentralized operators. The major evidence for centralized direction is the extreme correlation in content. This content includes domain targeting, large swaths of exact text matches, and commonly shared media. Many of these images and videos do not otherwise appear in reverse image searches, so they are likely produced centreally for this campaign.

But within this homogeneity of content, there is also substantial heterogeneity that speaks to decentralized operations. Within the large set of hundreds of accounts, there are smaller subsets of 5-15 accounts each that are technically much more similar to each other than they are to the other accounts in the campaign. The accounts in each “pod” were created on or near the same date. The profile/banner pictures show a pattern. The same client is used to post to Twitter. They post content in the same order. In a technical section, below, we formalize the detections of these pods, but the variation among the pods coupled with the homogeneity within indicates that the distribution of this content was not centralized.

The combination of the exact character of the content (original media and verbatim tweets) and the structure of the operation combines to tell an intriguing story about how this network is managed. To easily allow operators with different technical setups, including using several desktop and several mobile device types, to access timely specialized content, it must be hosted online in a centralized repository. But since the dominant mode of posting is from Twitter from Android, this repository is probably optimized for posting from mobile to Twitter. Finally, as none of the novel media are found on other social media platforms, the campaign (and likely the software for distributing the centralized content) is likely focused exclusively on Twitter.

Figure 19: Example of cartoons shared by pro-Indian government network.

Pro-Indian Network #2: Four Coordinated Accounts Identified

Finally, when we examine co-linker accounts promoting particular hashtags, we find one hashtag that stands out to an extreme degree: #GlobalThreatCCP. Co-linking accounts are those that post the same link within a very small time window - such as 60 seconds - to increase traffic and amplify the link's message. 'Colink' degree measures the total number of times a Twitter account (i.e., node in the network) has shared the same URL within the specific time window with its neighbors. We can interpret this as nodes with higher co-link degree are more strongly coordinating. There are only a few co-linker accounts promoting this hashtag, but they are doing so at extremely high rates.

23

When we examine these accounts we find that, like the previous, much larger network, they are pro-Indian government and extremely critical of China and Pakistan. Unlike the previous network, however, these accounts do not purport to be individuals from India. They have taken the persona of non-Indians, mostly Westerners (using stolen profile images) but also one Pakistani. These accounts all operate for the same few hours a day and all employ Twitter Web App exclusively. They routinely attempt to engage with particular accounts, to replying politically engaged accounts (journalists, professors, etc.) with pro-Indian content. These accounts mentioned AUKUS in tweets targeting individuals including, for example, Roberta Metsola (President of the European Parliament), Chris Horton (Taiwan-based Western journalist), and Diana Johnson (British Parliament).

hashtags

	convo	co_linker	tweet_count	user_count	short_day
0	#PEACE	0.967391	92	4	03/15/2023
0	#AUKUSnPOTUS	0.968421	95	4	03/15/2023
0	#JournalismIsNotACrime	0.968421	95	4	03/15/2023
0	#FreeSpeech	0.954315	197	10	03/15/2023
0	#GlobalThreatCCP	1.000000	98	3	03/15/2023
0	#NuclearWeaponsn3	1.000000	58	1	03/14/2023
1	#FreePress	0.953333	150	8	03/15/2023
1	#JobsNotWar	1.000000	44	1	03/16/2023
2	#HealthcareNotWarfare	1.000000	44	1	03/16/2023
2	#GlobalThreatCCP	1.000000	92	1	03/19/2023
3	#FreeJulianAssangeNOW	0.942308	52	4	03/15/2023
3	#NuclearWeaponsn3	0.989247	93	2	03/17/2023
3	#FreeJulianAssange	0.917197	157	14	03/15/2023

Figure 20: A list of hashtags that are being promoted by accounts that often post the same link within a very small time window, with #GlobalThreatCCP highlighted.

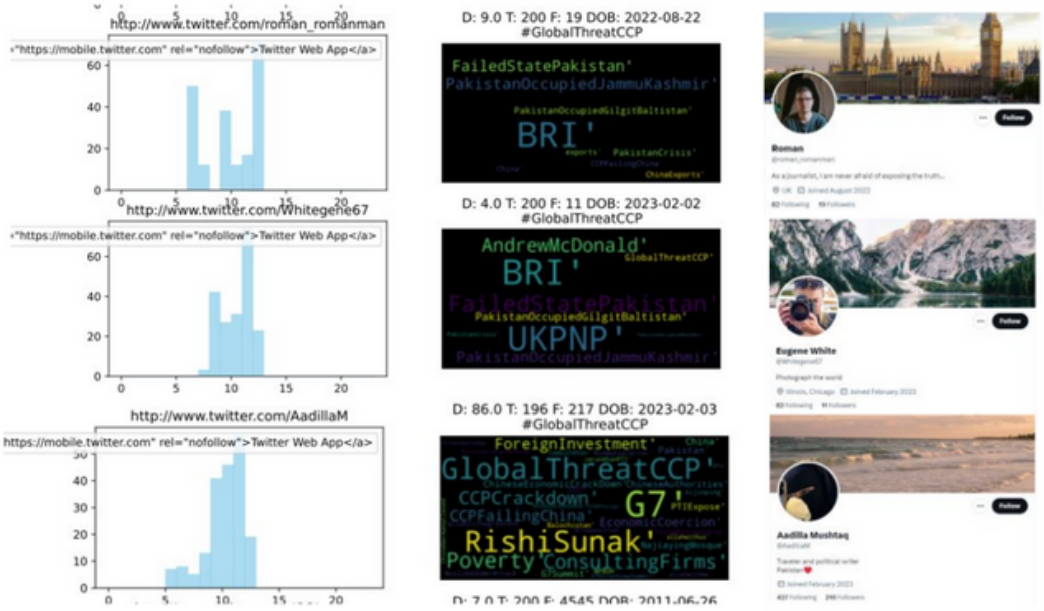

Figure 21: Profile, histogram, and word cloud of three accounts associated with the pro-Indian cluster.

Figure 22: Example tweet from the Pro-Indian cluster.

Detecting accounts likely controlled by a single operator

We use an advanced algorithm to detect multiple coordinating accounts that are likely to be controlled by a single operator or source. While coordination is a useful signal in the workflow for detecting inauthentic accounts, it tells us little about how many operators are controlling the accounts. Here we use sophisticated sequence analysis techniques and information theory to find clusters of accounts that we term 'pods', that is, multiple accounts controlled by the same human operator or source (e.g., controlled by computer code).

Our algorithm provides a degree of flexibility for the detection process. It is not restricted to detecting only identical messages posted in sequence. It uses information theoretic techniques to find messages that are similar but may differ slightly. For example, messages may be modified programmatically to add different @mentions of other accounts before or after the message. In this case, the message is still mostly the same but may have a few words (@mentions of accounts) that are different.

Likewise, messages may not be posted in perfect sequential order each time. For example, a human operator may become distracted and miss a browser tab when cycling through the accounts they control to post messages on. Similarly, a human operator in a content farm controlling multiple physical devices in a room might move left to right clicking on each one, then simply move back through from right to left. In this way, our approach quantifies the degree of 'sequentiality' in clusters of coordinating accounts, enabling the analyst to focus attention on clusters that have abnormally low metrics for sequential activity (where a low score means a high degree of sequential behavior).

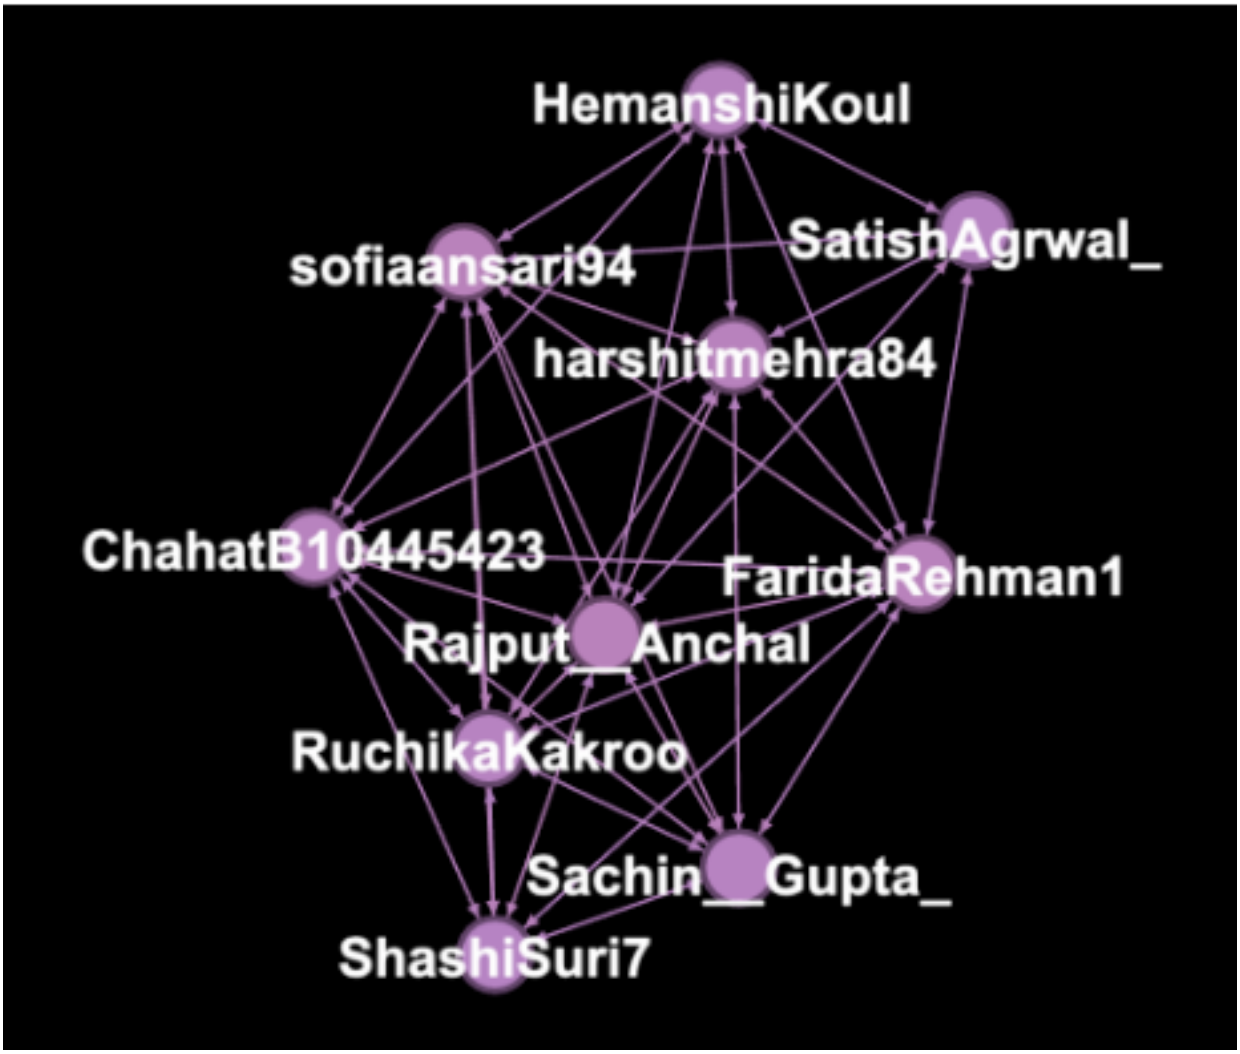

message_0	The Indian Navy, which is undergoing a modernisation process, may get yet another boost as France has offered a major nuclear submarine deal to India.&
message_1	-US-led alliances slowly but surely encircling China
-AUKUS alliance reveals its nuclear plans while emerging US-Japan-Philippines military alignment
message_2	-Australia Says Nuclear Subs Needed to Counter Militarization
-PM Albanese called AUKUS deal “the biggest single investment in Australia&#8
message_3	Eyeing China, Joe Biden and allies unveil nuclear-powered submarine plan for Australia
https://t.co/J1yfnwGLnd
message_4	China poses an “epoch-defining systemic challenge” to the U.K. and its allies, British Prime Minister Rishi Sunak said Sunday, as the U.K. g

Figure 23: Network visualization of Indian Pod # 1 and the five messages sent by this pod.

The first five messages sent by Indian Pod #1 are listed below:

1. The Indian Navy, which is undergoing a modernisation process, may get yet another boost as France has offered a major nuclear submarine deal to India. As part of the deal, France will become a part of India's program to develop 6 nuclear submarines.....
2. US-led alliances slowly but surely encircling China. AUKUS alliance reveals its nuclear plans while emerging US-Japan-Philippines military alignment changes the calculus around Taiwan
3. Australia Says Nuclear Subs Needed to Counter Militarization. PM Albanese called AUKUS deal "the biggest single investment in Australia's defense capability in all of our history.
4. Eyeing China, Joe Biden and allies unveil nuclear-powered submarine plan for Australia.
5. China poses an "epoch-defining systemic challenge" to the UK & its allies, British PM Sunak said, as the UK govt would spend an extra $6B on its nuclear-armed submarine fleet & replenishing munitions stockpiles to bolster support for Ukraine & deter China

Figure **X** shows the first pod of interest in the AUKUS dataset, along with the five messages they posted. This is a pod of ten coordinating accounts that post five messages in the exact same order each time. These accounts have since been suspended by Twitter / X. The accounts are pushing a narrative that underscores the growing concerns and defensive postures of various nations against China. This could be an attempt to shape public opinion, create a sense of urgency around defense, or legitimize defense spending and alliances.

n particular, the tweets consistently emphasize different nations' concerns about China, whether it's the US, Australia, or the UK. By highlighting these concerns, the accounts may be attempting to portray China as a significant and common threat. Likewise, the tweet about France offering a major nuclear submarine deal to India might be an attempt to cheer-lead for India's growing defense partnerships and its importance on the global stage.

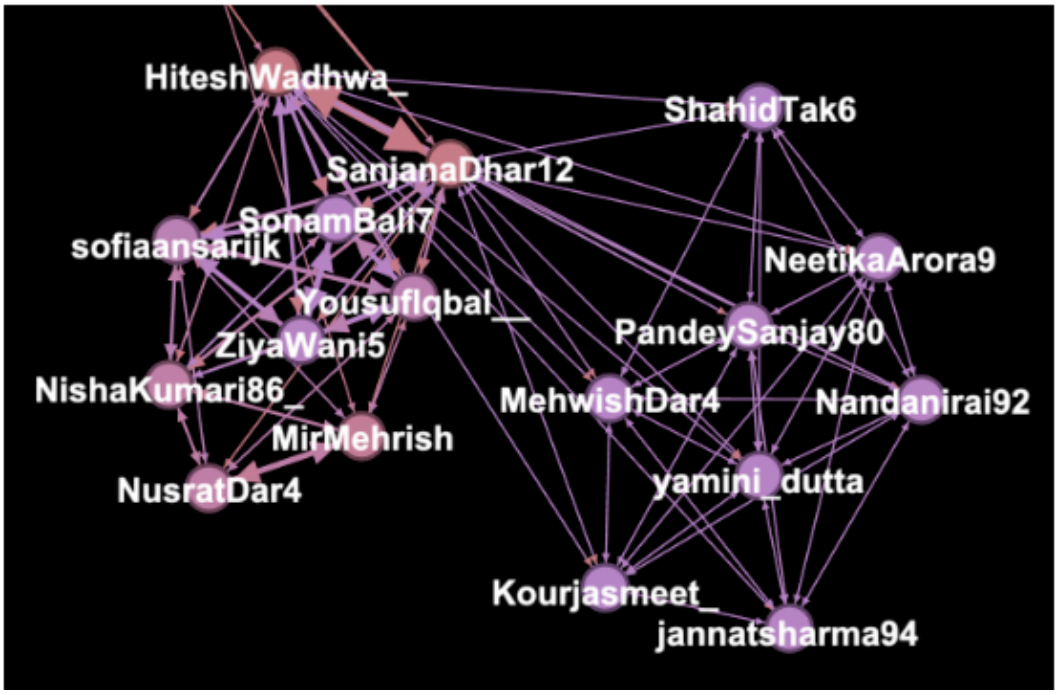

Figure 24: Network visualisation of Indian Pod # 2.

Figure **X2** provides another example of an Indian pod. This pod has interesting structural characteristics because it consists of two densely connected sub-pods. These sub-pods are connected by two accounts: Hiteshwadhwa and SanjanaDhar12. These two accounts are, in turn, connected to a third pod which we discuss further below. What is interesting about this pod is that the messaging is not always in perfect sequential order. The colors of the nodes represent the degree to which the accounts are sequentially aligned in their messaging with their local neighborhood. The color gradient is from orange to purple, that is, from least to most sequentially aligned. We observe that the right-hand pod is very sequentially aligned while the left-hand pod has a little more disorder in the messaging patterns, as shown by the pinkish hue of the nodes on the left-hand side compared to the purplish hue on the right. We hypothesize that accounts in this pod may be swapped in and out of different operators, suggesting that two to three operators control them. We also find that the right-hand pod accounts are all suspended, however three of the nine left-hand side accounts are not suspended. This suggests that Twitter / X is using different approaches to detecting inauthentic accounts given that all of these accounts are clearly part of the same campaign and probably controlled by a very small set of operators.

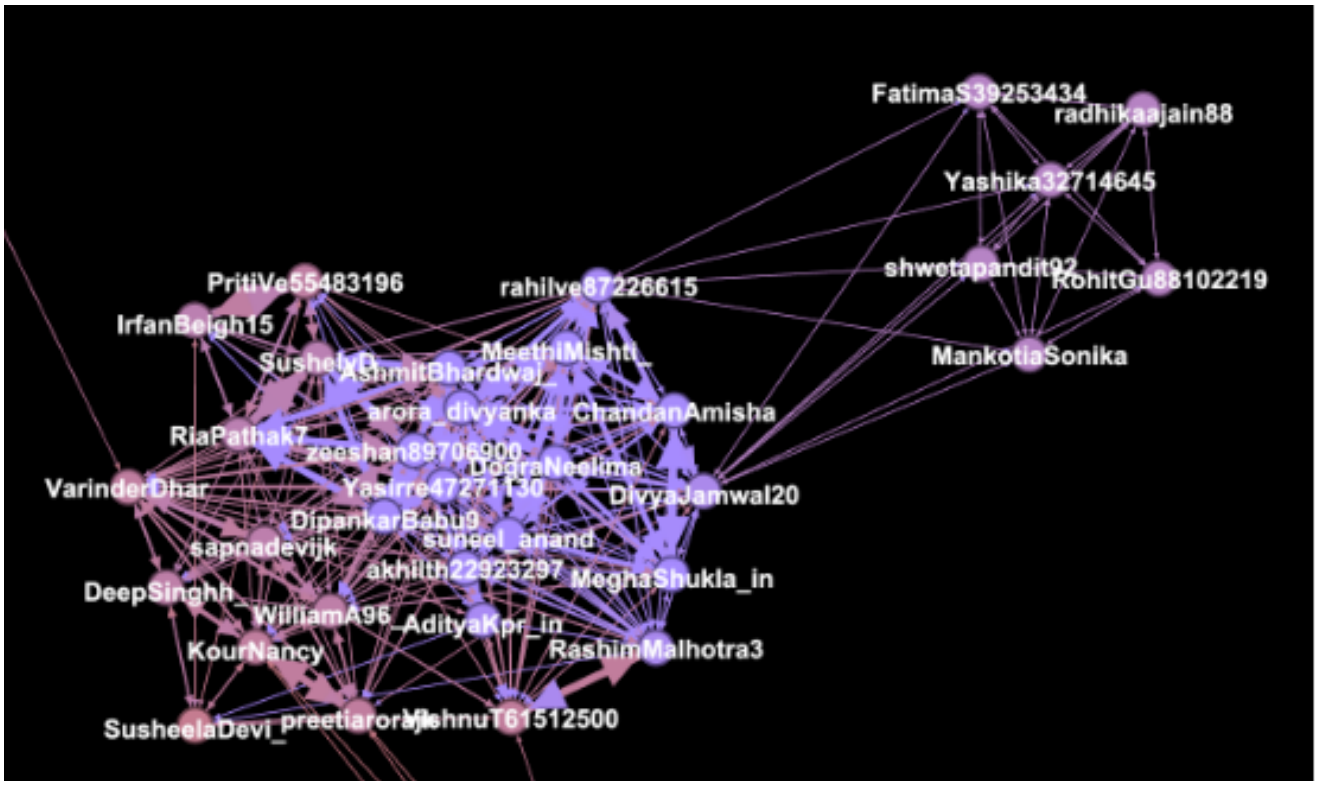

Figure 25: Network visualization of Indian pod # 3.

Finally, we focus attention on the third Indian pod as shown in Figure X3. Similar to Pod #2 above, this network consists of two sub-pods as shown on the left- and right-hand sides. However, the left-hand pod is analytically interesting because even though all the accounts are strongly coordinated and thus densely clustered, not all of them are posting sequentially.

29

Discussion and Implications

The analytic processes demonstrated in this analysis of Twitter discourse addressing AUKUS can give similar insight to a range of topics across communities and platforms. Our processes are an efficient mechanism for giving context and meaning to large volumes of data, especially about the sets of organized and coordinated accounts that are pushing select objects of influence. By integrating the actor-centric approach from Clemson and the time-centric approach from QUT, we provide a fully integrated system for influence campaign detection that provides world class accuracy and sensitivity. As we bring this technology to market, we expect this integration to even further deepen.

Further, we know that putting these processes within a software package offering the ability to find, manipulate, and explore data through a user dashboard can further enhance the power of these tools. A user dashboard would not only add speed and efficiency, they would also make these tools accessible to analysts with less training and start-up time.

There are currently several tools available on the market for graphing social media networks and understanding the flow of information. No tool, however, gives the depth of context we have developed the ability to offer. And certainly no tool allows analysts to make judgements about coordination and authenticity as quickly and with this degree of accuracy. We developed these tools because we needed them and they were not available elsewhere. We would like to change that.

30

Appendix

Appendix A: Details on Methodological Approaches

QUT Analytic Process

The QUT team utilized the Coordinated Detection Toolkit to statistically determine the extent of disinformation methods by examining temporally coordinated behavior (Graham, 2020). Coordinated behavior occurs when multiple accounts are posting or reposting the same content, repeatedly and within a short time frame of each other. Examples of coordinated online behaviors include:

1. Co-retweet: reposting the same post as another account
2. Co-tweet: tweeting identical text to another account
3. Co-linking: posting the same link as another account
4. Co-reply: replying to the same post as another account

The QUT analytical process also permits multi-behaviour analysis. This allows us to consider accounts and actors that are engaged in more than one of the above behaviors; for example, a user who is co-replying and co-tweeting. Our coordination analysis approach also allows for sequential analysis; that is, a set of accounts that engage in one of the above coordinated behaviors one after another, in a sequence. Sequential coordinated activity is a further indicator that the activity is inauthentic and likely connected to an information operation campaign. For example, bot accounts will often be controlled by a source script that has a list of accounts that are iterated over at regular intervals, thereby resulting in the accounts posting one after the other every time.

The importance of our ability to track multi-behaviour coordination cannot be overstated; in real world case studies, it is likely that inauthentic accounts in coordinated campaigns are behaving inauthentically in multiple ways simultaneously and are not limiting themselves to simply co-replying, for example. Using the AUKUS debate dataset, we constructed a Multi-behaviour Coordination Network consisting of four behavior types and thus four edge types: co-retweet, co-tweet, co-reply, and co-link. We set a time window of 60 seconds and minimum edge weight of 2, meaning that accounts (nodes) need to have performed the same action within 60 seconds of each other at least twice (edge weight >= 2) in order to be included in the Multi-behaviour Coordination Network. The network consists of 6,087 nodes and 199,093 edges.

31

Clemson Analytic Process

While the QUT approach focuses on identification of coordination through sequential behavior, the Clemson team takes a different, but complementary, approach. The Clemson process first takes full data sets on a given topic (e.g. “AUKUS” or “Xinjiang” or “Presidential Debate”) and identifies all of the accounts within that conversation. Each account is then given one or more flags if that account has an attribute that makes it suspicious. These might include the “baby” flag if the account is newly created, the “bot” flag if it posts at an extreme rate, the “first” flag if it posts frequently in the first second of the minute, or the “egg” flag if it has the default profile. No one flag is a definitive marker of inauthenticity (everyone’s account was once brand new, after all). Flags merely indicate accounts of interest and how they are interesting.

Next, the Clemson process identifies within a topic data set objects of potential influence. Actors may work to either promote or demote the relative prominence of a variety of objects of influence, including hashtags, domains, specific accounts, or a given narrative within a topic. The process then examines what accounts are sharing messages regarding each of these objects of influence and what the relative mix of flags for each object may be. Normally this is done separately for each day appearing in the dataset.

Any dataset will have a distribution of account flags messaging about objects of influence. When an unusual mix of flags is messaging about a given object of influence on a given day (e.g., a high proportion of baby accounts talking about a particular hashtag on a particular day), however, this will warrant further exploration.

Appendix B: Context on Organizations Behind Networks

1. Bharatiya Janata Party Bot Networks

India is an under-discussed actor in literature surrounding bots, trolls, and information operations. The strategic use of social media in the nation begins from the top down; Indian Prime Minister Narendra Modi is highly active on social media and one of the most followed world leaders (Sharma, 2023). Modi, as leader of the Bharatiya Janata Party (BJP), utilized Twitter campaigns to great effect in the 2014 and 2019 general elections (Sharma, 2023). Indian society has a unique social structure that includes a wide mix of religions, castes, languages, regions, and classes; each area of tension provides fuel for bots and trolls (Biju & Gayathri, 2023).

Political actors across nations and the political spectrum are utilizing social media for disinformation and partisan messaging, with India being no exception. The BJP Information Technology Cell is a department of the Indian government responsible for the management of social media campaigns for the party and its members. Previous white papers from think tanks have published findings of extensive networks of fake local media outlets with inauthentic accounts supporting Indian interests from 65 countries (Machado et al., 2020).

2. Miles Guo Network

Miles Guo – variously known as Guo Wengui, Miles Kwok, or Ho Wan Kwok – is a significant but complicated figure in online misinformation networks. In 2014, Guo fled his native China for New York due to accusations of corruption and various other crimes from Chinese officials (Hilgers, 2018). Upon moving to the United States, Miles Guo ingratiated himself with some of the most prominent Republicans and became a business associate of Steve Bannon (Osnos, 2023). Known for his extensive criticism of the Chinese Communist Party (CCP), Guo took his messaging further and, in 2020, announced via live stream that he and Bannon had established the New Federal State of China, a shadow government (Osnos, 2022).

Guo has ties to Trump political strategist Steve Bannon and together they have been criticized for masterminding a sprawling network of online disinformation that spreads conspiracy theories, anti-CCP sentiment, anti-vaccination rhetoric, and false allegations of election fraud in the 2020 US Presidential election. Previous work utilizing the same methodological techniques as this case study found a coordinated bot-like network funded by Miles Guo, self-styled as the

'Whistleblower Movement' or 'Whistleblower Revolution' (爆料革命). This network was found to be spreading anti-Chinese Communist Party and COVID-19 conspiracy theory content in a study of Australian political disinformation on Twitter (Graham & FitzGerald, 2023). The network of bot-like accounts all had almost identical profiles, were posting in both English and Chinese, and specifically amplifying right-wing accounts in the Australian Twittersphere (Graham & FitzGerald, 2023). Guo's coordinated information operation has been previously reported on by the Australian Strategic Policy Institute (ASPI), when it spread state-aligned disinformation around protests in Hong Kong (Uren et al., 2019).

References

Biju, P. R., & Gayathri, O. (2023). Self-breeding Fake News: Bots and Artificial Intelligence Perpetuate Social Polarization in India's Conflict Zones. The International Journal of Information, Diversity, & Inclusion (IJIDI), 7(1/2), Article 1/2. https://doi.org/10.33137/ijidi.v7i1/2.39409

Graham, T., & FitzGerald, K. M. (2023). Politicians Spreading Disinformation on Social Media: An Exploratory of Craig Kelly MP on Twitter. Media International Australia, Under Review.

Hurst, D. (2021, October 28). Scrapping submarines deal broke trust, Macron tells Australian PM. The Guardian. https://www.theguardian.com/world/2021/oct/28/france-seeks-tangible-actions-from-australia-after-submarines-row

Kahn, L. (2023, June 12). AUKUS Explained: How Will the Trilateral Pact Shape Indo-Pacific Security? Council on Foreign Relations. https://www.cfr.org/in-brief/aukus-explained-how-will-trilateral-pact-shape-indo-pacific-security

Machado, G., Alaphilippe, A., Adamczyk, R., & Gregoire, A. (2020). Indian Chronicles: Deep dive into a 15-year operation targeting the EU and UN to serve Indian interests. EU DisinfoLab. https://www.disinfo.eu/publications/indian-chronicles-deep-dive-into-a-15-year-operation-targeting-the-eu-and-un-to-serve-indian-interests/

Murphy, K., & Hurst, D. (2021, September 16). Australia nuclear submarine deal: Aukus defence pact with US and UK means $90bn contract with France will be scrapped. The Guardian. https://www.theguardian.com/australia-news/2021/sep/16/australia-nuclear-submarine-deal-contract-france-scrapped-defence-pact-us-uk

Osnos, E. (2022, October 17). How a Tycoon Linked to Chinese Intelligence Became a Darling of Trump Republicans. The New Yorker. https://www.newyorker.com/magazine/2022/10/24/how-a-tycoon-linked-to-chinese-intelligence-became-a-darling-of-trump-republicans

Osnos, E. (2023, March 16). What Secrets Does the "Donald Trump of Beijing" Know? | The New Yorker. The New Yorker. https://www.newyorker.com/news/daily-comment/what-secrets-does-the-donald-trump-of-beijing-know

Pengelly, M., & Hawkins, A. (2023, March 15). Chinese business tycoon and Bannon ally Guo Wengui arrested in $1bn fraud conspiracy. The Guardian. https://www.theguardian.com/world/2023/mar/15/ho-wan-kwok-arrest-fraud-conspiracy-steve-bannon

Sharma, N. (2023). Populism and social media use: Comparing the Indian Prime Minister Narendra Modi's strategic use of Twitter during the 2014 and the 2019 election campaigns. Media Asia, 50(2), 181–203. https://doi.org/10.1080/01296612.2022.2135269

U.S. Embassy in Canberra. (2023, March 13). AUKUS Joint Leaders' Statement. U.S. Embassy & Consulates in Australia. https://au.usembassy.gov/aukus-joint-leaders-statement/