Networks and Influencers in Online Propaganda Events: A Comparative Study of Three Cases in India

ANONYMOUS AUTHOR(S)*

The structure and mechanics of organized outreach around certain issues, such as in propaganda networks, is constantly evolving on social media. We collected Twitter data on two propaganda events and one non-propaganda event with varying degrees of organized messaging that emphasize a one-sided narrative on complex events, and perform a comparative analysis of the user and network characteristics of social media networks around these events. Our study reveals clearly distinguishable traits across events. We find that influential entities like prominent politicians, digital influencers, and mainstream media prefer to engage more with social media events with lesser degree of propaganda while avoiding events with high degree of propaganda, which are mostly sustained by lesser known but dedicated *micro-influencers*. We also fintwork communities of events with high degree of propaganda are significantly centralised with respect to the influence exercised by their leaders.

CCS Concepts: • **Applied computing** → *Computers in other domains*; • **Information systems** → **Information retrieval**; **Social networks**.

Additional Key Words and Phrases: propaganda, misinformation, social media analysis, social network, influencers

ACM Reference Format:
Anonymous Author(s). 2018. Networks and Influencers in Online Propaganda Events: A Comparative Study of Three Cases in India. In . ACM, New York, NY, USA, 19 pages. https://doi.org/XXXXXXX.XXXXXXX

1 INTRODUCTION

Social media is one of the most important sources of news and a range of daily information in today's age. Alongside changing peoples' access to networks as well as giving them voice and communities to engage with, social media has also been a central force in the spread of misinformation and propaganda [74] with speed and virality that was rarely possible in the era of institutionally-managed broadcast media. Examples abound where a social media movement leads to popularization of a propagandist, and often fringe idea with ferocity and footprint in a way that not only obscures reality, but actively makes the empirically verifiable knowledge in an issue look like minority perspectives.

Much has been studied on the provenance and propagation of information [46], on networks [14], on the role of individual actors in it [21], and on the affective nature of content [61] and its likelihood of reaction. In short, we know by now that bad behavior or content arousing negative emotions gets more attention online [10]. However, there are still gaps in understanding what brings these factors together, specifically, how particular events and emotive content that are driven by vested interests differ from those that do not share the same attributes. For this, we turn to the notion of propaganda, which encompasses content engineered and driven by vested interests.

The Institute for Propaganda Analysis defines propaganda as the *"expression of opinion or action by individuals or groups deliberately designed to influence opinions or actions of other individuals or groups with reference to predetermined ends"*[17]. Propaganda varies in its degree and type depending on the event or topic considered, and can have both

Manuscript submitted to ACM

positive and negative valence, based on whether its purpose is to degrade or to edify [40].While propaganda may either be crafted to positively project the image of those that drive it, in this study, we mainly focus on negative propaganda, i.e., messaging that aims to negatively project an ostensible antagonist entity or group. Such entities have been theorized in the political science literature as an 'enemy' within a broader political spectacle [31], one whose discrediting or construction as problematic is fundamental to the legitimacy of the actor that does the propagandist work. Our focus in this work is propaganda in short-form messaging as seen on Twitter, specifically that with identifiable negative framing aimed to discredit one group (please refer to table 6 for some examples).

We present a comparative analysis of the social media user and network characteristics around two propaganda events and one non-propaganda event, each with different target groups involved. Based on qualitative analysis, we term an event a *propaganda event* if there exists a significant presence of propagandist tweets, which attempt to influence the opinions of Twitter users to the predetermined perspective of one sender, whether individual or group. The events considered in this study carry varying degrees of propagandist messaging. The research question that we want to address through this work is: *Does the social media output on Twitter around different types of events with varying degrees of propagandist messaging leave different traces in their user and network characteristics?*

We conducted this study in India, the world's largest democracy, which has the highest installed base for several social media platforms including WhatsApp, Instagram, YouTube, and the third-highest installed base for Twitter. India has also been in the news over social media use in large part due to news coverage of polarization and political media use. Our analysis is solely on Twitter, which despite relatively lower installed base than Facebook and WhatsApp has grown to be the preferred platform of outreach for politicians and journalists, playing a central role in driving the national discourse [39]. We examine three events that saw significant online activity in this paper – the conversations around CAA-NRC, the Indian National Budget of 2022, and the Boycott Bollywood movement online (hereafter referred to as CAA, Budget, and Bollywood).

The Citizenship (Amendment) Act (CAA) was passed by the Indian Parliament in December 2019, with the purpose of providing a pathway to Indian citizenship for persecuted religious minorities, other than Muslims, from other countries, who arrived in India before the end of December 2014. The NRC stands for the National Register of Citizens, an effort to gather proof of residency and claim to citizenship of people living within India's borders. The two were meant to work in tandem. It was a marquee act of the right wing Hindu nationalist BJP government, which came to power on the plank of 'Hindutva' or a notion of statecraft and polity with Hindu culture and tradition at its center. As the BJP's first explicit act of codifying faith as a factor for access to citizenship, in a nation-state that has secularism enshrined in its constitution, the law was widely discussed in the media in India and abroad, with sharp polarization on its consequences on the idea of an inclusive India. With its stringent requirements for proving citizenship, the act also set off alarm bells for citizens who had no connection to neighboring states, since proving citizenship involved significant documentation [66], and the onus of proving origin was much deeper on Muslims.

The second event considered that saw widespread social media activity was the Union Budget Report for the year 2022. The budget is one of the most important policy documents in India, since it lays out annually what are the funds apportioned to various ministries and initiatives, as well as the rates of taxation for various goods and services. The budget is typically a politically contentious issue since it inevitably has winners and losers. In the 2022 Budget, the government primarily pushed for capital expenditure and private investment, but drew flak from the opposition and social media for its lack of initiatives to enhance employment opportunities or mitigate inflation [54], which has been a significant issue in the aftermath of COVID-19.

The third event which had a significant social media footprint was the *Boycott Bollywood* campaign that frequently encourages people to boycott films made by a certain cohort of the Hindi-language film actors industry, referred to as *Bollywood* since it is largely run from the state of *Mumbai*, earlier known as *Bombay*. The campaign started gaining momentum starting in the early days of COVID-19, with the death of a popular Indian actor named *Sushant Singh Rajput (SSR)* in June 2020 [3]. His fans claimed that the actor's demise was either a pre-meditated murder or a suicide triggered by systemic bullying from a nepotistic clique in the Mumbai film industry. Although initial investigations into the matter concluded that the death was by suicide, the campaign to investigate various angles as well as to boycott the output of the industry was kept alive by a set of dedicated social media fans or micro-influencers [26]. While the actor's death was a precipitating event, Boycott Bollywood took a communal turn as the discourse turned its attention to Muslim actors in the Hindi film industry [53], some of who are the highest-paid and most successful stars.

We classify the CAA and Boycott Bollywood as propaganda events, while the Budget as a non-propaganda event. There are two distinguishing features that help frame this categorization. First, unlike the Budget which is a yearly regular cycle event, the CAA and Boycott Bollywood are unilateral actions that elicit reactions from those opposed to it. Although it is also true that the Budget typically has detractors from the opposition, this tends to be around line items of expenditures, which are inherently political, but the idea of the budget itself is part of the normal business of governance. Second, with both the CAA and with Boycott Bollywood, there is an explicit agenda for one party to propose a normative set of reasons to propagate its position, and outwardly demonstrate support for it. The Budget, while an important annual issue, is typically delivered as a missive with its justifications at the point of delivery, but with very little effort to get citizen buy-in thereafter. While both CAA and Boycott Bollywood are propaganda events – in the former, the state is much more directly involved in the propaganda, making it a lot closer to a traditional state propaganda event in which a government takes the initiative to sell its position on an issue. On the other hand, in the latter, a mix of private and public actors, presenting themselves as upholding a national interest, take part in the propaganda activity. It is also important that a state propaganda event has access to the institutions and means of outreach that can reinforce a narrative with an air of sanctioned authority [64].

To understand the distinguishing characteristics of the three events, we perform analysis of their user and retweet network characteristics. Our analysis reveals that: (A) The social media discourse of a political non-propaganda event like Budget is primarily led by mainstream media houses and accounts of prominent politicians (those with a high number of followers beyond a certain threshold), while that of a political propaganda event (CAA) is led by media houses, digital influencers, and low influence political accounts, (B) Political propaganda events (CAA) also see the active involvement of lesser known micro-influencers, (C) Propaganda events with chiefly non-political target users (Boycott Bollywood) see minimal participation from media, influencer, and political accounts, but are sustained almost solely by micro-influencers, (D) Political accounts choose to engage on an event's discourse on Twitter based on their stature – non-propaganda events show more engagement by prominent politicians and vice versa, and (E) Propaganda communities show significantly higher inequality in their retweet networks than non-propaganda communities, with respect to their influence distribution, i.e., a few influential leaders guide the narrative of the entire community for the propaganda events.

While several studies have been carried out in the area of propaganda detection on Twitter, to the best of our knowledge, ours is the first work that performs a comparative analysis of the distinguishing user and network characteristics of events with varying levels of propaganda. Thus, we show that not only do propaganda events vary from non-propaganda events in terms of these characteristics, but propaganda events even show variation among themselves, depending on their type. Additionally, we also present interesting findings around temporal evolution of event communities. The framework of analysis proposed can act as the first step towards early modeling of the degree of propagandist messaging for a social media event based on its type and target groups, which in turn affect its user and network characteristics.

2 RELATED WORK

Several earlier studies have touched upon the topic of propaganda analysis. Some of them apply machine learning based approaches for the analysis of textual content of messages using supervised techniques to predict propaganda, in a multi-class [33, 35, 63] and binary classification set up [8, 9, 56], and unsupervised techniques to detect propaganda [25, 76] or to identify groups of users spreading propaganda [51, 57]. Many of these approaches use word n-grams, formal text representations like readability and writing style, background information [44], and lexicon based methods for propaganda prediction [43, 72]. Some studies [23, 27, 55] focus on developing extractive propaganda detection methods using deep neural networks, which unlike the other studies, detect the technique of propaganda spread along with the propagandist text fragment or span. In terms of the type of content studied for propaganda analysis, previous research focuses on both textual content [4, 47, 68] and multimodal content present in tweets [28, 37, 65, 67].

Initial network analysis techniques for propaganda detection focus on the problem of node detection, i.e., these techniques classify a user node as propagandist or not in isolation [41], using supervised techniques on various user level features including network features [75, 77], textual content used [62], and profile information [5, 38, 42, 71]. Approaches on the role of media accounts [6] and bots in propaganda spread on Twitter have also been developed [1, 2, 11, 13, 16, 73, 75]. Additionally, there are other works that target propaganda analysis on social media using both content and network analysis [18, 34, 36].

There are three limitations of these approaches: (A) Many of them significantly rely on annotated datasets/corpora used for training, thus being dependent on the amount of annotation, its subjectivity, and its quality, (B) They are vulnerable to automated text generation techniques that are often used to game them, and (C) It is often difficult to tag a node/user as propagandist in isolation, since such nodes often deliberately use techniques to evade detection. Researchers in this area thus eventually focused on detecting coordinated behavior [39] among a set of user nodes, rather than considering a node in isolation, using semi-supervised and unsupervised techniques [15, 30]. In this direction, studies on detecting suspicious user connectivity patterns [20, 45, 58] and temporal tweeting/retweeting patterns [19, 29, 49] have been developed.

Our work is motivated by these previous studies, and attempts to look at propaganda analysis through a different lens. Instead of focusing on detection/classification of propaganda or propagandist users, we study how different events consisting of some level of propagandist messaging vary from each other with respect to their user and network characteristics. More specifically, we want to understand the various domains to which the primary users driving the social media discourse on propaganda and non-propaganda events belong, and the way their influence is distributed in the event discourse, over time.

3 METHODOLOGY

3.1 Data Sampling

For each event, we collected publicly available Twitter data using the Twitter V2 API. These included basic metadata (such as the time of posting, contents, and mentioned accounts) as well as network level features (such as user IDs referenced in the tweets and the number of retweets received). For each event, we first collected a set of tweets using

Event	Initial Hashtag(s) (H1)	All Hashtags (H2)
CAA-NRC	caanrc	caanrc, caanrcnpr, anticaa, caanrcprotest
Budget 2022	budget2022, unionbudget	budget2022, unionbudget2022, budget22, unionbudget22, unionbudget
Boycott Bollywood	BoycottBollywood	BoycottBollywood, Boycottbollywoodcompletely, Boycottbollywoodforever, BoycottAlia, BoycottGangubai

Table 1. Hashtags used for tweet data collection for the three events: H1 contains the initially selected most popular hashtags, while H2 contains the relevant hashtags frequently co-occurring with H1

manually selected exemplar hashtags (H_1) which are listed in Table 1. These hashtags were discussed and shortlisted based on their frequency and contextual relevance to the events concerned. Next, we found the most frequently co-occurring hashtags (from the set of 200 most retweeted tweets with the hashtags in $H1$) which were not included in H_1 for each event, and used them to expand the hashtag/keyword set for each event to H_2 (also listed in table 1). In doing so, we manually inspected every new hashtag and only added them to H_2 if they were unambiguously event-specific (for instance, the generic keyword 'protest' was not used to collect CAA-NRC specific tweets, even though the event was marked by long standing offline protests). This process weights precision over recall. We then used the expanded keyword/ hashtag set H_2 to collect all the tweets over the respective event timelines (mentioned below). The objective behind this two-step data collection process was to optimise for recall, i.e., to ensure that our datasets adequately represent the universe of tweets for the events. We finally obtained 104533, 843936, and 643663 tweets for CAA-NRC (Dec-April 2019), Budget 2022 (Jan-June 2022), and Boycott Bollywood (Jan-April 2022) respectively, including retweets. This data was used for all further analysis in this study. The full list of hashtags (sets H_1 and H_2) used are presented in Table 1.

3.2 Annotation Method

To understand the degree of propagandist messaging for each event, tweets corresponding to the events had to be categorized into one of the two categories of *Propaganda* (P) and *Non-Propaganda* (NP). Three annotators who were not part of the group of authors of this study manually annotated the top 500[1] most retweeted messages corresponding to each event into one of the two categories. As part of the manual coding process, the annotators first had a group discussion and created a coding scheme with descriptions for inclusion and exclusion criteria (table 6). The group discussions included conversations on the news and metaphor around the events, which were important since there is significant use of innuendo and contextual information matters. Our coding scheme contained a set of example tweets for each event, their gold standard labels (P/NP), and an explanation of why they belonged to one of the two categories.

The labeling of tweets was done based on the presence of 18 commonly known features of propagandist messaging as elaborated in existing literature [48]. We considered tweets as propagandist when (A) They contained at least one of these 18 features as part of their textual content and (B) They directly vilified or discredited a particular group or individual with an explicit agenda to demonstrate support towards a predetermined belief or movement.

Next, to evaluate the annotators' understanding of the exercise, they were presented a set of 20 tweets (10 propaganda and 10 non-propaganda tweets) and asked to categorize them. This exercise revealed an agreement of above 90% for each event among the annotators, i.e., for 10% of the tweets, at least one annotation differed from the others. For these tweets, the annotators decided on a gold standard annotation along with the rationale behind it after discussion.

[1]We considered the top 500 retweets based on the cumulative distribution of retweets received for each event.

Event	Primary Target Groups	Frequency	Propaganda (% of tweets)	#unique authors
CAA-NRC	Religious groups, political entities	One time (since 2019)	42%	447
Budget 2022	Political entities	Yearly	11%	459
Boycott Bollywood	Hindi-language film fraternity (actors, directors, and producers)	Sustained movement since 2020	98%	184

Table 2. Categorization of events based on qualitative analysis: Three annotators independently annotated 500 top retweeted tweets for each event as propaganda/non-propaganda

Finally, we provided the annotators a set of 500 most retweeted tweets for each event as the main annotation task. Post annotation, to measure the inter-annotator agreement, we calculated the Cohen's Kappa statistic which provided a K value of 0.9, indicating a significantly high degree of agreement. For the tweets carrying disagreement in annotation, we considered the majority label as the final annotation. Based on the percentage of annotated propagandist tweets among these 500 tweets, we present the categorization of the three events in table 2. From this qualitative analysis, we find that Budget is essentially a non-propaganda event, since it contains very few propagandist tweets (11%) among its top retweeted messages. The highest propagandist messaging is seen for Boycott Bollywood (98%) followed by CAA (42%), both of which are propaganda events as also established in previous studies [24, 52]. Arguably, the sampling of Boycott Bollywood lends closer to propaganda, because it already uses the term *Boycott* (e.g., as opposed to just using *Bolywood* which would have some share of pro- and anti-messaging). Whereas the CAA has propaganda messaging by virtue of being a highly emotive, hot-button issue.

We annotated the top 500 most retweeted messages for the events, since the top retweeted messages provide us a fairly accurate idea of what is "successful" as far as the outreach goes, and in turn, defines the nature of dominant Twitter discourse around the event communities considered. If a significant fraction of the most retweeted messages around an event are propagandist in nature, we can consider the event a propaganda event as these messages carry the maximum impact on the community.

The percentage of propagandist messages in this cohort of 500 most retweeted tweets provides us a few hints towards the difference in the nature of the two propaganda events: First, we find CAA to be a much balanced event when compared to Boycott Bollywood, which has minimal number of non-propagandist tweets, again, likely because of their respective semantic construction. CAA is a nationally prominent offline event with very significant press coverage, as well as strong opinions from social media users from both sides of the sprectrum – those with a predetermined agenda of discrediting a religious group (propagandist) and those with situational updates, opinions, and positive messages about the ongoing protests (non-propagandist). On the other hand, Boycott Bollywood is a primarily social media movement, whose very purpose is to discredit the Hindi-film fraternity, thus by construction itself propagandist and aiming to vilify film-stars and Bollywood in general. Second, we find significantly lesser number of unique authors for Boycott Bollywood's most retweeted tweets, when compared to CAA. This is indicative of a small group of people driving the entire event's narrative through a dedicated social media presence.

3.3 Community Detection

We wanted to understand how the various user accounts on Twitter cluster together to form retweet communities, the characteristics of these users, and how these communities evolve over time for each event. Each community within an event is a set of user accounts who heavily retweet each other, while engaging in minimal cross-community retweeting.

Event	#Nodes	#Edges	Modularity
CAA-NRC (P)	10655	17898	0.58
Budget 2022 (NP)	3631	13033	0.59
Boycott Bollywood (P)	11948	124299	0.27

Table 3. Details of the retweet networks formed by the core set of users for the events (P: Propaganda, NP: Non-Propaganda): Budget shows the minimum modularity, indicative of significant cross-community retweeting

For this purpose, we first considered the retweet network for accounts belonging to mainstream media houses (identified from the *DISMISS* dataset[7]), major digital influencers[2] (identified from the *DISMISS* dataset[7]), political accounts (identified from the *NivaDuck* dataest [59]), above 80th percentile user accounts with highest retweets received, and the above 99th percentile user accounts with highest number of retweets posted on the event based on the cumulative distributions (CDF) of the corresponding metrics. This sampling was done to ensure that we remove the inconsequential users from the network, who were only sporadically active during the event timeline, and also to reduce the time complexity of the community analysis exercise. This set of users finally considered for our analysis is referred to as the *core set* of users or *core users*. It is important to note that other than political, media, and influencer accounts, the core sets for each event consists of users who are influential in their own networks (even if they might not be well known outside their networks). We call these users the *micro-influencers* [60].

The retweet network for an event is a directed, unweighted graph with the core sets of user accounts as nodes, and directed edges (u, v) where an edge from node u to node v denotes that u has retweeted v at least once regarding the event, during the event timeline. The details of the retweet networks for each event after this sampling is shown in table 3. For community analysis on each event's retweet network, we used the Greedy Modularity community detection technique [22]. This maximization technique begins with each node in its own community, and repeatedly joins the pair of communities that lead to the largest modularity (a measure of isolation between communities), till convergence. We see in figure 1 the retweet network communities for the three events. It is clear from the figures that the event networks form around three to five major communities each. We also cross-checked this finding by calculating the number of intra-community edges for each community per event. For all further experiments on community analysis, we experimentally considered the top five communities in terms of their number of intra-community edges, since the rest consisted of significantly fewer number of retweet edges (< 100), and these outliers were hence considered inconsequential.

4 RESULTS

4.1 User Characteristics

We find significant social media participation of political, media, and influencer accounts for the events considered, alongside sustained activity by lesser known micro-influencers in the two propaganda events (figure 6 in Appendix). In this section, we take a deeper look at these user characteristics empirically.

4.1.1 Media and Influencer Accounts. To understand the tweeting pattern of mainstream digital influencer and media accounts for the three events, we compare in figure 2 the number of overall tweets during the event timeline, with the number of tweets posted on event specific hashtags, by digital influencers and mainstream media accounts. From the

[2]Influencers are individuals or accounts who command a large following on social media and wield influence either directly or through their ability to get second-order engagement in their extended social media networks.

(a) CAA-NRC (Propaganda) (b) Budget 2022 (Non-Propaganda) (c) Boycott Bollywood (Propaganda)

Fig. 1. Plots of the retweet-networks for the three events reveals the presence of 3-5 major communities in each event (Each color represents a separate community). Plots have been made with the Gephi software, using OpenOrd as a layout algorithm. Community outliers with fewer than 100 retweet edges have been excluded

plots, we see that both influencers and media houses tweet similarly across events. Budget (non-propaganda) shows a sharp peak and a subsequent drop in tweeting during the first week of February 2022 for both. This is the period when the Budget was presented in the Parliament, and most discussions on Twitter happened on this mainstream political event. CAA being a comparatively longer term political propaganda event shows a sharp peak for both media and influencers at the time of its inception (December 2019), but a gradual drop. Finally, Boycott Bollywood being a fringe movement with the highest propaganda, shows no tweeting by either the media or the influencers. This indicates that its social media campaign is primarily maintained by lesser known dedicated users or micro-influencers, and the mainstrem media/influencer accounts abstain from tweeting on it.

4.1.2 Political Accounts. We observed that the three events considered in this study show potential of politicians' interaction on social media. CAA being a national movement, involved long term protests for and against the decision of the ruling dispensation related to citizenship. Budget is a mainstream political event. Boycott Bollywood, while involving discussions primarily on Indian film fraternity, also brought in tweets by influential politicians [32]. Hence, we wanted to understand the participation of political accounts around these events. For this purpose, we plotted the number of original tweets posted by political accounts for the three events against the percentile of their number of followers in figure 3. The number of followers serves as an indicator of a politician's influence on social media.

We find that each event shows a different characteristic in terms of the political accounts interacting on them. For CAA and Boycott Bollywood (propaganda events), we see that the low influence political accounts (primarily in the 0-40 percentile range in terms of number of social media followers) interact the most on social media. Additionally, between the two events, CAA shows a significantly higher number of tweets by political accounts compared to Bollywood. This is due to the fact that CAA being a core political issue attracted a lot of tweets and propaganda from aggressive

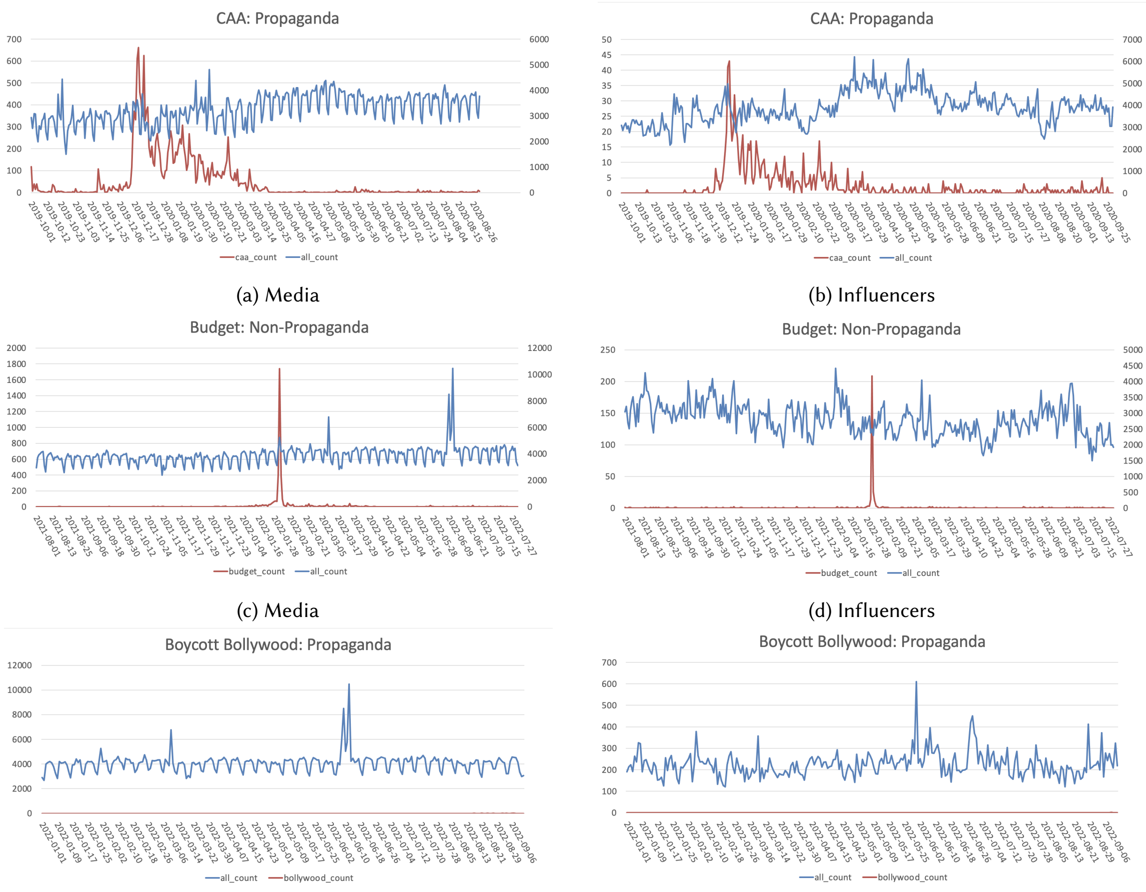

(a) Media (b) Influencers

(c) Media (d) Influencers

(e) Media (f) Influencers

Fig. 2. Number of original tweets posted by media houses (left) and digital influencers (right) on the events (in red) vs. number of all tweets posted by them irrespective of the event (in blue).

politicians with relatively lower number of followers [70]. Similarly, Boycott Bollywood being a fringe movement with mostly non-political target groups attracted fewer propagandist tweets mostly from low influence politicians. For both of these events, the high influence, prominent politicians stayed away from tweeting, probably owing to the necessity of maintaining their social and political image. Contrarily for Budget (non-propaganda), we find the activity is highly skewed towards the high influence politicians (80-100 percentile). Most of these political accounts tweeted on the justification or criticism of the mainstream political event.

4.2 Network Characteristics

In this section, we analyze the characteristics of the retweet network of core set of users (refer to section 3.3) in terms of the information spread and the community characteristics.

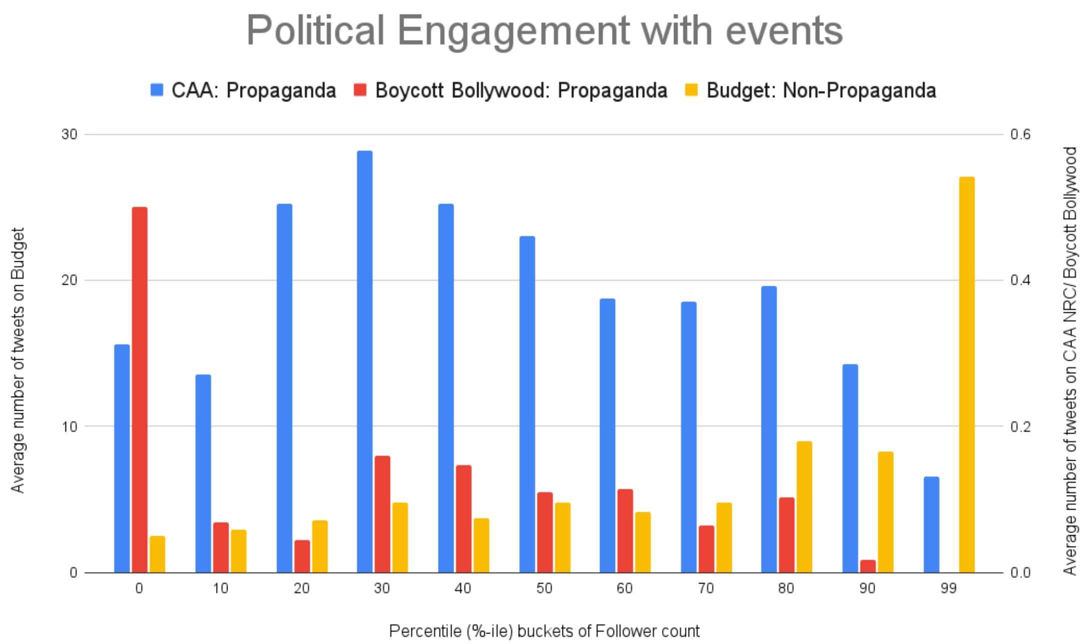

Fig. 3. Original tweets by political accounts against their number of followers: Prominent politicians are seen to abstain from tweeting on the propaganda events, while low influence politicians tweet significantly on them

4.2.1 Spread of Information Generated. We first analyze the spread achieved by tweets generated by these accounts within the entire event network (core and non-core users combined). This is done to understand the actual impact these core users have in the overall event narrative. We measure the spread achieved by a set of user accounts as the percentage of nodes in the event data who retweeted any original tweet by a member account in the set, during the event timeline. Table 4 shows the spread achieved by the different types of core users for each event.

Spreader	CAA-NRC	BoycottBollywood	Budget
Influencers	0.28	0.05	0.21
Politicians	1.03	0.03	0.52
Media houses	0.34	0.03	0.06
All (core set)	1.98	0.92	0.84

Table 4. Spread achieved by different types of core users (in %) for each event: The spread of a set of accounts represents the percentage of all user accounts in the event dataset who retweeted at least one tweet authored by any user in the set during the event timeline

We find that as expected, the two propaganda events CAA and Boycott Bollywood exhibit a significantly higher overall spread than Budget, a mainstream political, non-propaganda event. Additionally, Boycott Bollywood being a fringe propaganda movement, shows a much lower spread than CAA by media, political, and influencer accounts, since the former is chiefly sustained by lesser known micro-influencers. It is important to note that while the overall spread achieved by the core accounts for all three events is significantly low, our definition of spread only captures the 1-hop retweets received by the accounts. It will be interesting to measure the influence that these core accounts exercise, in terms of the multi-hop reach of their tweets in the event networks as part of our future work.

4.2.2 Community Analysis. As discussed in the Related Work section, propagandist users often work in coordination. A sign of such coordination is formation of multiple densely clustered user communities within the event network, which contain sub-groups of similar users in terms of their nature of discourse. We wanted to see if for the three events, the core user accounts form distinct communities with minimal inter-community communications. We find from table 3 that the CAA and Budget networks show a significantly higher modularity, when compared with the modularity of the Bollywood network. This is primarily due to the fact that Boycott Bollywood is a fringe propaganda movement where a set of lesser known (non-politician and non-influencer) but dedicated micro-influencers maintain the campaign and receive significant number of retweets from across multiple communities within the event. Thus, the communities are less isolated for the Boycott Bollywood event, compared to the other two.

Next, we analyze the temporal evolution of communities for each event. For this purpose, we consider the retweet network formed till the end of each month in the event timeline, since the start date of the event. Thus, the i^{th} retweet network (n_i) contains the retweet network of core users who participated till the i^{th} month since the start of the event. We considered monthly evolution of the retweet networks for ease of analysis, and also because we observed that it generally takes a month for significant change to happen in the networks in terms of the number of retweet edges. To understand how intra-community influence of accounts evolve over time, we next calculated the intra-community GINI coefficient for the top five communities for each event, considering the intra-community PageRank of the member user accounts. A GINI value closer to 0 indicates high equality in the distribution, while that closer to 1 indicates high inequality. Figure 4 shows the intra-community GINI coefficients over time for each event.

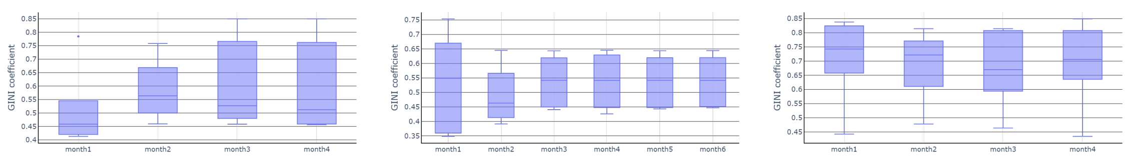

(a) CAA-NRC (Propaganda) (b) Budget 2022 (Non-Propaganda) (c) Boycott Bollywood (Propaganda)

Fig. 4. Intra-community GINI coefficient of the events over time: Each community within an event is identified using Greedy Modularity Maximization algorithm. For each community, we calculate the GINI coefficient of PageRank distribution among the member accounts.

Based on the median values, we find that Boycott Bollywood shows the maximum inequality in intra-community influence over time (for each pair of events, difference between median GINI values is significant with $p < 0.05$). From this finding, we can argue that Budget and CAA being political events with lesser degree of propaganda consist of communities in the retweet network that exhibit a more equitable distribution of intra-community influence. On the other hand, in Boycott Bollywood a set of highly influential leaders guide the propagandist narrative within their communities.

To also understand if the community leaders remain constant over time for the top five communities, we calculated the Jaccard Coefficient between the top 15 leaders (top three leaders from each of the top five communities) for each temporally consecutive pair of retweet networks. We find that initially for all three events, the leadership of the top communities show some change. However, over time the top communities stabilize in terms of leadership as seen from the high values of jaccard coefficient as shown in figure 5. The observable trait here is that CAA being a short term and

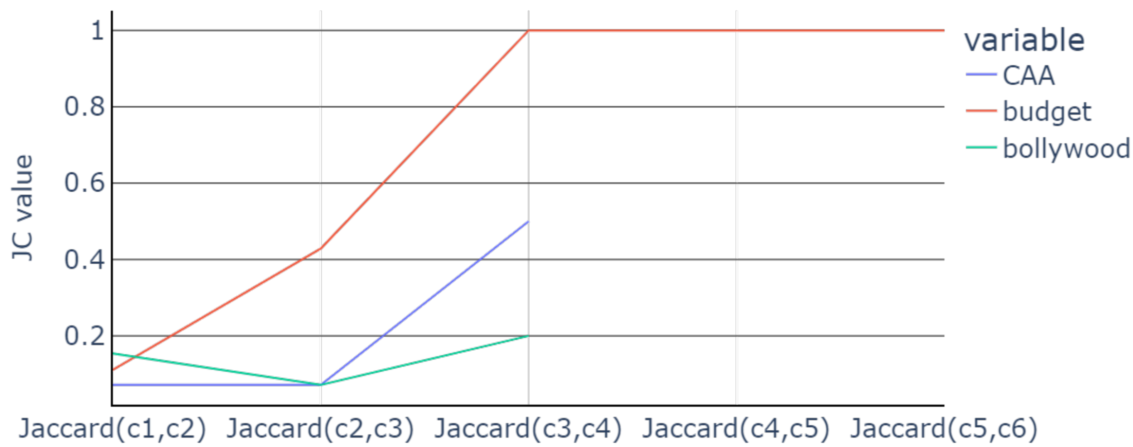

Fig. 5. Jaccard coefficient of top communities over time: ci denotes the set of top five communities for n_i, i.e., retweet network till the i_{th} month. $Jaccard(ci, cj)$ denotes the Jaccard coefficient between the top 15 leaders (top 3 leaders from each of the five communities) of sets ci and cj

one-time event shows a significant change in leadership of its top communities in the first couple of months, and shows a much higher slope in the plot compared to the other two events, which are either repetitive or have been sustained for a significant time. Thus, it seems that for political movements like CAA, the initial leadership is usurped eventually by newer leaders as the movement matures in the social media. Long sustained movements like Boycott Bollywood do not show an equally significant change in leadership (as also evident from figure 4), owing to the maturity of the movement. Finally, mainstream political events show the highest commonality in leadership , since only a small subset of politicians and experts capture the narrative eventually.

Finally, we see if the leaders of the top five communities are political accounts, media accounts, or digital influencers. Table 5 shows the percentage of each of these types of accounts among the community leaders over time. Corroborating our earlier findings, we see that Budget (non-propaganda) being a mainstream political event exhibits the maximum intra-community leadership by media houses and politicians, followed by CAA. Digital influencers constitute the community leadership more for the propaganda event CAA than Budget. Boycott Bollywood shows negligible presence of any of these types of entities, being a fringe propaganda event (the intra-community leaders being micro-influencers). Additionally, we find that the communities and their leadership for Budget becomes fixed the third month onwards, since there is negligible discussion after the third month on the event (as also seen from figure 4). On the contrary, CAA communities show significant change in community leadership till the fourth month of the event. CAA being a political propaganda event with public protests also exhibits some community leadership by lesser known micro-influencers, unlike Budget where the community leadership is captured exhaustively by one of political/media/influencer accounts (and the percentages add up to 100).

	CAA (P)	Budget (NP)	Bollywood (P)
month 1	(20,7,40)	(60,7,33)	(0,0,0)
month 2	(27,13,27)	(33,0,67)	(0,0,0)
month 3	(13,0,33)	(40,0,60)	(0,0,7)
month 4	(20,0,27)	(40,0,60)	(0,0,0)

Table 5. Percentage (rounded off) of media houses, influencers, and politicians in the community leadership over time, for the three events (P: Propaganda, NP: Non-Propaganda)

5 CONCLUSION AND DISCUSSION

The rise of social media as an important source of news and daily information has also resulted in it being a primary force in the spread of misinformation and propaganda. While several previous studies have tackled the problem of propaganda detection at the message, user, and group levels, we examine the contours of three Indian events – CAA-NRC, Budget 2022, and Boycott Bollywood – through a comparative analysis of their social media user and network characteristics, provided that these events are already known to consist of varying degrees of propaganda. CAA-NRC, an act passed by the Indian Parliament, drew global criticism since it codified faith as a factor for access to Indian citizenship. The Union Budget Report for 2022 is a politically contentious issue that exhibited polarised messaging on social media from both the ruling dispensation and its opposition. Boycott Bollywood is a social media campaign that encourages people to boycott films made by a certain group of the Hindi-language film industry, involving repetitive propagandist messaging from a group of dedicated campaigners. CAA-NRC is categorized as a political propaganda event, since it targets radical policy reforms at the national level, and contains significant propagandist messaging targeting a religious group. Boycott Bollywood is categorized as a fringe propaganda event as it attempts to target mostly the Hindi-film industry with highest levels of propaganda. Budget 2022 is categorized as a political non-propaganda event.

We find that the events show significant differences in their user and network level characteristics. CAA shows a significantly higher activity by digital influencers compared to the other two events, even after mainstream media attention has subsided for the issue. This is indicative of influencers' need to engage over mainstream political events to appeal to the masses. Neither mainstream media nor influencer accounts show significant activity for Boycott Bollywood, since it is a fringe movement almost solely maintained by lesser known but dedicated users or micro-influencers. For these two events, politicians with relatively lesser number of followers on Twitter show the most tweeting activity, probably to stay relevant in their political careers. On the other hand, prominent politicians tend to tweet less on these movements to avoid the risk of over-engaging in propaganda events, while showing significant activity for Budget, which is a non-propaganda event.

In terms of the retweet network characteristics, CAA and Budget show a higher modularity, or isolation of intra-network communities, compared to Boycott Bollywood. This is indicative of significant cross-community retweeting for the latter. We also find that these communities consist of "community leaders" – influential accounts who attract the most retweets and direct the event discourse – that become increasingly static as the communities evolve. For all three events, the intra-community influence distribution is somewhat skewed, i.e., the community leaders hold a disproportionately higher influence than others in terms of retweets received within the communities. However, this skew is the highest for the fringe propaganda event Boycott Bollywood. Thus, social media events with high degree of propaganda are seen to exhibit higher centralization of intra-community influence across the event timeline compared to others with lesser amount of propagandist messaging involved. Finally, a deeper look at these community leaders reveals that while both CAA and Budget shows significant community leadership by mainstream media and political accounts, the former also contains communities that are led by lesser known micro-influencers who do not belong to any of these cohorts.

Thus, we find that other than distinguishable traits between propaganda and non-propaganda events, there are traits that vary among propaganda events as well. These observations open different directions for future work. First, our analyses can be used to study other events across geographies, and it will be important to observe if these findings hold for them. Second, while we studied each of the events in isolation in this work, an interesting direction would be to understand if and how these event communities overlap or merge over time as shown in previous works [50]. Third, an attempt to understand the connection between textual content of tweets, with the user and network dynamics for the events can reveal more of the underlying causes for the patterns observed. In this direction, study of multi-lingual and code-mixed content [12, 69] will be indispensable. Finally, attempts can be made to model the degree of propagandist messaging among multiple social media events, using their distinguishable user and network characteristics.

6 LIMITATIONS

This paper presents a comparative analysis of the user and network characteristics of three Indian events with varying levels of propaganda. We find that different events show clearly distinguishable patterns in these characteristics over time, which could aid platform policies in regulating propagandist communities of practice. However, this work must be viewed as a preliminary attempt at comparatively characterising propagandist events and has some limitations. First, the analysis needs to be replicated across multiple other events with varied nature, frequency, and target groups to make the insights more generalized. Second, while we have tried to perform the qualitative analysis of tweets with minimal bias and using established qualitative analysis techniques, the definition of propaganda itself is subjective, and we have used the extant definition according to previous literature in the field, to only consider negative propaganda – propaganda that vilifies a group or individual. It will also be interesting to consider propagandist messaging carried out with a positive intent and understanding user and network characteristics around it. Finally, the user characteristics studied in this paper heavily depend on the existing datasets on Indian political, media, and influencer accounts. These datasets need to be updated with time to include newer entities for the completeness of analysis.

REFERENCES

[1] Norah Abokhodair, Daisy Yoo, and David W McDonald. 2015. Dissecting a social botnet: Growth, content and influence in Twitter. In *Proceedings of the 18th ACM conference on computer supported cooperative work & social computing*. 839–851.

[2] Nitin Agarwal, Samer Al-Khateeb, Rick Galeano, and Rebecca Goolsby. 2017. Examining the use of botnets and their evolution in propaganda dissemination. *Defence Strategic Communications* 2, 1 (2017), 87–112.

[3] Syeda Zainab Akbar, Ankur Sharma, Himani Negi, Anmol Panda, and Joyojeet Pal. 2020. Anatomy of a Rumour: Social media and the suicide of Sushant Singh Rajput. *arXiv preprint arXiv:2009.11744* (2020).

[4] Saja Aldera, Ahmad Emam, Muhammad Al-Qurishi, Majed Alrubaian, and Abdulrahman Alothaim. 2021. Online extremism detection in textual content: a systematic literature review. *IEEE Access* 9 (2021), 42384–42396.

[5] Amit A Amleshwaram, Narasimha Reddy, Sandeep Yadav, Guofei Gu, and Chao Yang. 2013. Cats: Characterizing automation of twitter spammers. In *2013 Fifth International Conference on Communication Systems and Networks (COMSNETS)*. IEEE, 1–10.

[6] Cynthia Andrews, Elodie Fichet, Yuwei Ding, Emma S Spiro, and Kate Starbird. 2016. Keeping up with the tweet-dashians: The impact of'official'accounts on online rumoring. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*. 452–465.

[7] Arshia Arya, Soham De, Dibyendu Mishra, Gazal Shekhawat, Ankur Sharma, Anmol Panda, Faisal Lalani, Parantak Singh, Ramaravind Kommiya Mothilal, Rynaa Grover, et al. 2022. DISMISS: Database of Indian Social Media Influencers on Twitter. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 16. 1201–1207.

[8] Alberto Barrón-Cedeno, Giovanni Da San Martino, Israa Jaradat, and Preslav Nakov. 2019. Proppy: A system to unmask propaganda in online news. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 9847–9848.

[9] Alberto Barrón-Cedeno, Israa Jaradat, Giovanni Da San Martino, and Preslav Nakov. 2019. Proppy: Organizing the news based on their propagandistic content. *Information Processing & Management* 56, 5 (2019), 1849–1864.

[10] Jonah Berger and Katherine L Milkman. 2012. What makes online content viral? *Journal of marketing research* 49, 2 (2012), 192–205.

[11] David M Beskow and Kathleen M Carley. 2018. Bot-hunter: a tiered approach to detecting & characterizing automated activity on twitter. In *Conference paper. SBP-BRiMS: International conference on social computing, behavioral-cultural modeling and prediction and behavior representation in modeling and simulation*, Vol. 3. 3.

[12] Shankar Biradar, Sunil Saumya, and Arun Chauhan. 2021. Hate or non-hate: Translation based hate speech identification in code-mixed hinglish data set. In *2021 IEEE International Conference on Big Data (Big Data)*. IEEE, 2470–2475.

[13] Gillian Bolsover and Philip Howard. 2019. Chinese computational propaganda: Automation, algorithms and the manipulation of information about Chinese politics on Twitter and Weibo. *Information, communication & society* 22, 14 (2019), 2063–2080.

[14] Alexandre Bovet and Hernán A Makse. 2019. Influence of fake news in Twitter during the 2016 US presidential election. *Nature communications* 10, 1 (2019), 1–14.

[15] Chiyu Cai, Linjing Li, and Daniel Zeng. 2017. Detecting social bots by jointly modeling deep behavior and content information. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*. 1995–1998.

[16] Guido Caldarelli, Rocco De Nicola, Fabio Del Vigna, Marinella Petrocchi, and Fabio Saracco. 2020. The role of bot squads in the political propaganda on Twitter. *Communications Physics* 3, 1 (2020), 1–15.

[17] Hadley Cantril. 1938. Propaganda analysis. *The English Journal* 27, 3 (1938), 217–221.

[18] Akemi Takeoka Chatfield, Christopher G Reddick, and Uuf Brajawidagda. 2015. Tweeting propaganda, radicalization and recruitment: Islamic state supporters multi-sided twitter networks. In *Proceedings of the 16th annual international conference on digital government research*. 239–249.

[19] Nikan Chavoshi, Hossein Hamooni, and Abdullah Mueen. 2016. Debot: Twitter bot detection via warped correlation.. In *Icdm*. 817–822.

[20] Aditya Chetan, Brihi Joshi, Hridoy Sankar Dutta, and Tanmoy Chakraborty. 2019. Corerank: Ranking to detect users involved in blackmarket-based collusive retweeting activities. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*. 330–338.

[21] Rajdipa Chowdhury, Sriram Srinivasan, and Lise Getoor. 2020. Joint estimation of user and publisher credibility for fake news detection. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. 1993–1996.

[22] Aaron Clauset, Mark EJ Newman, and Cristopher Moore. 2004. Finding community structure in very large networks. *Physical review E* 70, 6 (2004), 066111.

[23] Giovanni Da San Martino, Seunghak Yu, Alberto Barrón-Cedeno, Rostislav Petrov, and Preslav Nakov. 2019. Fine-grained analysis of propaganda in news article. In *Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (EMNLP-IJCNLP)*. 5636–5646.

[24] Saloni Dash, Arshia Arya, Sukhnidh Kaur, and Joyojeet Pal. 2022. Narrative Building in Propaganda Networks on Indian Twitter. In *14th ACM Web Science Conference 2022*. 239–244.

[25] Guillaume Daval-Frerot and Yannick Weis. 2020. WMD at SemEval-2020 Tasks 7 and 11: Assessing humor and propaganda using unsupervised data augmentation. In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*. 1865–1874.

[26] Sandipan Deb. August 2022. The politics of boycott Bollywood. https://www.moneycontrol.com/news/trends/entertainment/the-politics-of-boycott-bollywood-9058811.html.

[27] Daryna Dementieva, Igor Markov, and Alexander Panchenko. 2020. SkoltechNLP at SemEval-2020 Task 11: Exploring unsupervised text augmentation for propaganda detection. In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*. 1786–1792.

[28] Dimitar Dimitrov, Bishr Bin Ali, Shaden Shaar, Firoj Alam, Fabrizio Silvestri, Hamed Firooz, Preslav Nakov, and Giovanni Da San Martino. 2021. Detecting propaganda techniques in memes. *arXiv preprint arXiv:2109.08013* (2021).

[29] Andrej Duh, Marjan Slak Rupnik, and Dean Korošak. 2018. Collective behavior of social bots is encoded in their temporal twitter activity. *Big data* 6, 2 (2018), 113–123.

[30] Juan Echeverrï£¡ a, Emiliano De Cristofaro, Nicolas Kourtellis, Ilias Leontiadis, Gianluca Stringhini, and Shi Zhou. 2018. LOBO: Evaluation of generalization deficiencies in Twitter bot classifiers. In *Proceedings of the 34th annual computer security applications conference*. 137–146.

[31] Murray Edelman. 1988. *Constructing the political spectacle*. University of Chicago Press.

[32] TimesNow Entertainment Desk. August 2022. Akshay Kumar lands in trouble; BJP MP Subramanian Swamy issues legal notice, accuses team of distorting Ram Setu saga. https://www.timesnownews.com/entertainment-news/akshay-kumar-lands-in-trouble-bjp-mp-subramanian-swamy-issues-legal-notice-accuses-team-of-distorting-ram-setu-saga-bollywood-news-entertainment-news-article-93830877.

[33] Vlad Ermurachi and Daniela Gifu. 2020. UAIC1860 at SemEval-2020 Task 11: Detection of propaganda techniques in news articles. In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*. 1835–1840.

[34] Emilio Ferrara. 2017. Contagion dynamics of extremist propaganda in social networks. *Information Sciences* 418 (2017), 1–12.

[35] Dmitry Grigorev and Vladimir Ivanov. 2020. Inno at SemEval-2020 Task 11: Leveraging Pure Transformer for Multi-Class Propaganda Detection. *arXiv preprint arXiv:2008.11584* (2020).

[36] Stefano Guarino, Noemi Trino, Alessandro Celestini, Alessandro Chessa, and Gianni Riotta. 2020. Characterizing networks of propaganda on twitter: a case study. *Applied Network Science* 5, 1 (2020), 1–22.

[37] Mahdi Hashemi and Margeret Hall. 2019. Detecting and classifying online dark visual propaganda. *Image and Vision Computing* 89 (2019), 95–105.

[38] Maryam Heidari, James H Jones, and Ozlem Uzuner. 2020. Deep contextualized word embedding for text-based online user profiling to detect social bots on twitter. In *2020 International Conference on Data Mining Workshops (ICDMW)*. IEEE, 480–487.

[39] Brian Heredia, Joseph Prusa, and Taghi Khoshgoftaar. 2017. Exploring the effectiveness of twitter at polling the united states 2016 presidential election. In *2017 IEEE 3rd International Conference on Collaboration and Internet Computing (CIC)*. IEEE, 283–290.

[40] Garth S Jowett and Victoria O'donnell. 2018. *Propaganda & persuasion*. Sage publications.

[41] Sneha Kudugunta and Emilio Ferrara. 2018. Deep neural networks for bot detection. *Information Sciences* 467 (2018), 312–322.

[42] Sangho Lee and Jong Kim. 2014. Early filtering of ephemeral malicious accounts on Twitter. *Computer communications* 54 (2014), 48–57.

[43] Yingya Li, Jieke Zhang, and Bei Yu. 2017. An NLP analysis of exaggerated claims in science news. In *Proceedings of the 2017 EMNLP Workshop: Natural Language Processing meets Journalism*. 106–111.

[44] Ying Lin, Joe Hoover, Gwenyth Portillo-Wightman, Christina Park, Morteza Dehghani, and Heng Ji. 2018. Acquiring background knowledge to improve moral value prediction. In *2018 ieee/acm international conference on advances in social networks analysis and mining (asonam)*. IEEE, 552–559.

[45] Shenghua Liu, Bryan Hooi, and Christos Faloutsos. 2017. Holoscope: Topology-and-spike aware fraud detection. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*. 1539–1548.

[46] Yang Liu and Yi-Fang Wu. 2018. Early detection of fake news on social media through propagation path classification with recurrent and convolutional networks. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 32.

[47] Cristian Lumezanu, Nick Feamster, and Hans Klein. 2012. # bias: Measuring the tweeting behavior of propagandists. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 6. 210–217.

[48] Giovanni Da San Martino, Stefano Cresci, Alberto Barrón-Cedeño, Seunghak Yu, Roberto Di Pietro, and Preslav Nakov. 2020. A survey on computational propaganda detection. *arXiv preprint arXiv:2007.08024* (2020).

[49] Michele Mazza, Stefano Cresci, Marco Avvenuti, Walter Quattrociocchi, and Maurizio Tesconi. 2019. Rtbust: Exploiting temporal patterns for botnet detection on twitter. In *Proceedings of the 10th ACM conference on web science*. 183–192.

[50] Liz McQuillan, Erin McAweeney, Alicia Bargar, and Alex Ruch. 2020. Cultural convergence: Insights into the behavior of misinformation networks on twitter. *arXiv preprint arXiv:2007.03443* (2020).

[51] Bhaskar Mehta, Thomas Hofmann, and Peter Fankhauser. 2007. Lies and propaganda: detecting spam users in collaborative filtering. In *Proceedings of the 12th international conference on Intelligent user interfaces*. 14–21.

[52] Nivedita Menon. 2020. Hindu Rashtra and Bollywood: A New Front in the Battle for Cultural Hegemony. *South Asia Multidisciplinary Academic Journal* 24/25 (2020).

[53] Abir Misra. 2021. Can the Celebrity Speak? Controversies and the Eulogistic Fandom of Shah Rukh Khan. *American Behavioral Scientist* (2021), 00027642211042286.

[54] Shobhana K. Nair. February 2022. Union Budget 2022 reactions | Budget has no vision for employment generation: Opposition. https://www.thehindu.com/business/budget/union-budget-2022-opposition-reaction/article38359128.ece.

[55] Leonardo Nizzoli, Marco Avvenuti, Stefano Cresci, and Maurizio Tesconi. 2019. Extremist propaganda tweet classification with deep learning in realistic scenarios. In *Proceedings of the 10th ACM Conference on Web Science*. 203–204.

[56] Vitaliia-Anna Oliinyk, Victoria Vysotska, Yevhen Burov, Khrystyna Mykich, and Vítor Basto Fernandes. 2020. Propaganda Detection in Text Data Based on NLP and Machine Learning.. In *MoMLeT+ DS*. 132–144.

[57] Michael Orlov and Marina Litvak. 2019. Using behavior and text analysis to detect propagandists and misinformers on twitter. In *Annual International Symposium on Information Management and Big Data*. Springer, 67–74.

[58] Diogo Pacheco, Alessandro Flammini, and Filippo Menczer. 2020. Unveiling coordinated groups behind white helmets disinformation. In *Companion proceedings of the web conference 2020*. 611–616.

[59] Anmol Panda, A'ndre Gonawela, Sreangsu Acharyya, Dibyendu Mishra, Mugdha Mohapatra, Ramgopal Chandrasekaran, and Joyojeet Pal. 2020. Nivaduck-a scalable pipeline to build a database of political twitter handles for india and the united states. In *International Conference on Social Media and Society*. 200–209.

[60] Jiwoon Park, Ji Min Lee, Vikki Yiqi Xiong, Felix Septianto, and Yuri Seo. 2021. David and Goliath: when and why micro-influencers are more persuasive than mega-influencers. *Journal of Advertising* 50, 5 (2021), 584–602.

[61] Piotr Przybyla. 2020. Capturing the style of fake news. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 34. 490–497.

[62] Francisco Rangel and Paolo Rosso. 2019. Overview of the 7th author profiling task at PAN 2019: bots and gender profiling in twitter. In *Proceedings of the CEUR Workshop, Lugano, Switzerland*. 1–36.

[63] Hannah Rashkin, Eunsol Choi, Jin Yea Jang, Svitlana Volkova, and Yejin Choi. 2017. Truth of varying shades: Analyzing language in fake news and political fact-checking. In *Proceedings of the 2017 conference on empirical methods in natural language processing*. 2931–2937.

[64] Nelson Ribeiro. 2014. Broadcasting to the Portuguese Empire in Africa: Salazar's singular broadcasting policy. *Critical Arts* 28, 6 (2014), 920–937.

[65] Hyunjin Seo. 2014. Visual propaganda in the age of social media: An empirical analysis of Twitter images during the 2012 Israeli–Hamas conflict. *Visual Communication Quarterly* 21, 3 (2014), 150–161.

[66] Tathagat Sharma. 2020. The Pan India Experiment of NRC: History, Problems Associated and Lessons to Learn from Assam. (2020).

[67] Junfeng Tian, Min Gui, Chenliang Li, Ming Yan, and Wenming Xiao. 2021. MinD at SemEval-2021 Task 6: Propaganda Detection using Transfer Learning and Multimodal Fusion. In *Proceedings of the 15th International Workshop on Semantic Evaluation (SemEval-2021)*. 1082–1087.

[68] Javier Torregrosa, Joshua Thorburn, Raúl Lara-Cabrera, David Camacho, and Humberto M Trujillo. 2020. Linguistic analysis of pro-isis users on twitter. *Behavioral Sciences of Terrorism and Political Aggression* 12, 3 (2020), 171–185.

[69] Andrea Tundis, Gaurav Mukherjee, and Max Mühlhäuser. 2020. Mixed-code text analysis for the detection of online hidden propaganda. In *Proceedings of the 15th International Conference on Availability, Reliability and Security*. 1–7.

[70] Shylaja Varma. February 2020. "They Are Calling Me Terrorist": Defiant, BJP's Kapil Mishra Tweets Again. https://www.ndtv.com/delhi-news/delhi-violence-bjps-kapil-mishra-tweets-again-says-they-are-calling-me-terrorist-2185804.

[71] Nishan Chathuranga Wickramarathna, Thiruni D Jayasiriwardena, Malith Wijesekara, Pasindu Bawantha Munasinghe, and Gamage Upeksha Ganegoda. 2020. A framework to detect twitter platform manipulation and computational propaganda. In *2020 20th International Conference on Advances in ICT for Emerging Regions (ICTer)*. IEEE, 214–219.

[72] Alex Williams Kirkpatrick. 2021. The spread of fake science: Lexical concreteness, proximity, misinformation sharing, and the moderating role of subjective knowledge. *Public Understanding of Science* 30, 1 (2021), 55–74.

[73] William Williamson III and James Scrofani. 2019. Trends in detection and characterization of propaganda bots. (2019).

[74] Kai-Cheng Yang, Francesco Pierri, Pik-Mai Hui, David Axelrod, Christopher Torres-Lugo, John Bryden, and Filippo Menczer. 2021. The covid-19 infodemic: Twitter versus facebook. *Big Data & Society* 8, 1 (2021), 20539517211013861.

[75] Kai-Cheng Yang, Onur Varol, Clayton A Davis, Emilio Ferrara, Alessandro Flammini, and Filippo Menczer. 2019. Arming the public with artificial intelligence to counter social bots. *Human Behavior and Emerging Technologies* 1, 1 (2019), 48–61.

[76] Zichao Yang, Zhiting Hu, Chris Dyer, Eric P Xing, and Taylor Berg-Kirkpatrick. 2018. Unsupervised text style transfer using language models as discriminators. *Advances in Neural Information Processing Systems* 31 (2018).

[77] Zhi Yang, Jilong Xue, Xiaoyong Yang, Xiao Wang, and Yafei Dai. 2015. VoteTrust: Leveraging friend invitation graph to defend against social network sybils. *IEEE Transactions on dependable and secure computing* 13, 4 (2015), 488–501.

A USERS WITH CONSISTENT ACTIVITY

To understand the users who are consistently active over time through tweeting about the events, we consider users who have authored at least one original tweet about the event every month. We then generate monthly word-clouds of these user account handles based on their tweet count and color the consistent users in red. This allows us to visualise the cohesiveness and sustenance of discourse, in terms of which types of users are consistent in activity. In Figure 6, we show these clouds for the first three months of each event. We find that for CAA-NRC and especially for #BoycottBollywood, the discourse is sustained by individual accounts, many of whom do not belong to mainstream media, influencer, and political accounts. These users are often micro-influencers and have the social capital to sustain the discourse. We also note that highly propagandist fringe events, such as #BoycottBollywood are driven by an extremely consistent group of these micro-influencers, and have very limited activity from non-consistent (non-red) users. In essence, this is a distinguishing characteristic of such events. On the other hand, we observe that discussions on Budget 2022 are primarily sustained by media organizations that typically cover finance or business related news. The activity of these consistent accounts varies in intensity (number of tweets) over time for the events. These findings motivate us to further look into the participation of different types of user accounts in maintaining the discourse on the events.

B EXAMPLES OF TWEETS

Table 6 shows a snapshot of the coding scheme used to label tweets into the two categories of propaganda (P) and non-propaganda (NP) the three events. The example tweets also provide an idea of the kind of propaganda spread for the two propaganda events. While CAA-NRC shows propagandist tweeting primarily targeting a certain religious group, Boycott Bollywood consists of propagandist messaging around Sushant Singh Rajput's death, alongside encouraging people to abandon Hindi films made by a certain cohort of actors/directors.

Received 20 February 2007; revised 12 March 2009; accepted 5 June 2009

(a) CAA NRC: Month 1

(b) CAA NRC: Month 2

(c) CAA NRC: Month 3

(d) Budget: Month 1

(e) Budget: Month 2

(f) Budget: Month 3

(g) Boycott Bollywood: Month 1

(h) Boycott Bollywood: Month 2

(i) Boycott Bollywood: Month 3

Fig. 6. Word-clouds visualising consistently active accounts (in red) across events

Event (Label)	Example Tweets	Explanation
CAA-NRC (P)	"Show me one percent Muslim in India who has condemned Islamic radicalism. One percent Muslim who has condemned the AntiCAA protest. One percent Muslim who have condemned the rants of the Owaisi brothers. https://t.co/lkbAI2W3YJ"	The tweet negatively targets a religious group and carries unverified information, **exaggeration**, and **whataboutism**.
CAA-NRC (NP)	"Assam Cong MLA: Assam Govt should accept updated NRC, declare it as valid Read @ANI Story | https://aninews.in/news/national/politics/assam-cong-mla-assam-govt-should-accept-updated-nrc-declare-it-as-valid20230101085651/ #NRC #AssamGovernment #HimantaBiswaSarma #AssamCongress"	The news story supports a particular opinion that does not vilify any group or individual.
Budget 2022 (P)	"Rahul Gandhi claimed there are 2 Indias. Yes, there was an India before PM Sri @NarendraModi; India after Modi. India after Modi has achieved unprecedented prosperity of its people than what it was under Rule of Dynasts of Congress Speech in Parliament on #Budget2022 https://t.co/W9Z21SPPoD"	The tweet is emotive in nature and contains unverified claims, **exaggeration**, and **name calling**.
Budget 2022 (NP)	"On 1st February 2022 #Budget2022 was released in the houses of the parliament by respected Union Finance Minister #NirmalaSitharaman. Read the News to know what this year's budget has for the common people. https://t.co/tI2b6iAyQj"	The tweet shares a media article, stating the commencement of an event, and does not target any group or individual negatively.
Boycott Bollywood (P)	"SSR Sacrificed his life to Expose Gutterwood, Charsiwood Bullyweed #BoycottBollywood Sushant Conquered BWood https://t.co/UzHVLIpIrS"	The tweet contains negatively targets a group, and contains **exaggeration** and **name calling**.
Boycott Bollywood (NP)	"A 2015 Story is viral, where #ShahRukhKhan helped a #KashmiriPandit family in need of financial aid. A filmmaker, who is currently trending #BoycottBollywood in every news debate, then thanked #SRK for his big ... @srk never claimed such good deeds, even I have seen him doing so https://t.co/mXR9K1etgt"	The tweet narrates an incident and praises an individual. No trace of vilificationof any group or individual present.

Table 6. Coding scheme: Examples of propagandist and non-propagandist tweets identified using qualitative analysis (P: Propaganda, NP: Non-Propaganda). The tweets are annotated based on various attributes of propaganda [48] (in bold in the last column).