Journal of Information and Data Management, 2026, 17:1, doi: 10.5753/jidm.2026.7525
This work is licensed under a Creative Commons Attribution 4.0 International License.

Coordinated Behavior on Far-Right WhatsApp Groups in 2022 Brazilian General Elections

Fernanda Nascimento [**Universidade Federal do Ceará** | ***fernanda.nascimento@alu.ufc.br***]
José Maria Monteiro [**Universidade Federal do Ceará** | ***jose.monteiro@lsbd.ufc.br***]
Javam Machado [**Universidade Federal do Ceará** | ***javam.machado@lsbd.ufc.br***]

Computer Science Department, Universidade Federal do Ceará, Campus do Pici, Bloco 910, Pici, Fortaleza, CE, 60.440-900, Brazil.

Received: 30 January 2026 • **Published:** 27 November 2026

Abstract

This study investigates coordinated behavior within public WhatsApp groups associated with the Brazilian far-right during the 2022 general elections. Employing an unsupervised, network-based methodology grounded in Rapid Spread Networks, we operationalize coordination as the occurrence of highly similar messages (cosine similarity ≥ 0.70) disseminated by distinct users within a 60-second time window, and model the resulting interactions as weighted coordination networks for the pre-election (August–October 2022) and post-election (November 2022–February 2023) periods. Applied to a dataset of over 536,000 messages from 568 public groups, the used pipeline yields two sparse but structurally informative networks. The post-election network is larger (237 vs. 170 users; 176 vs. 129 edges) and exhibits substantially higher coordination intensity: the average edge weight more than doubles ($\bar{w} = 11.82$ vs. 5.59), and the fraction of strong connections ($w \geq 5$) grows from 13.2% to 23.3%. Core–periphery decomposition reveals a pronounced structural intensification: the proportion of users in the innermost coordination core expands from 2.35% in the pre-election period to 10.97% post-election, and among users active in both periods, the share occupying the strict core rises from 8.3% to 19.4%. Community detection (Louvain; $Q = 0.779$ and $Q = 0.745$, respectively) uncovers a dual dynamic: a proliferation of peripheral coordination dyads alongside the consolidation of larger, denser teams, the three largest post-election communities reaching 30, 29, and 7 members. Geospatial inference suggests a strategic geographic reorientation after the election, with the coordination core shifting from a multi-regional distribution centered on São Paulo (9 core users) to a marked concentration in Paraíba (17 core users). Semantic characterization confirms that the detected communities circulate content consistent with far-right electoral contestation and calls for institutional intervention, with misinformation scores remaining heterogeneous across communities despite a slight decrease in the global weighted average (0.349 to 0.313). These findings indicate that far-right coordination on WhatsApp does not dissipate after an electoral event but reorganizes into a more persistent and structurally concentrated form, with implications for the understanding of post-election influence operations on encrypted messaging platforms.

Keywords: Coordinated behavior; Information spreading; Misinformation; WhatsApp

1 Introduction

Recently, a significant political and electoral rise of the far-right has been observed in Brazil and various parts of the world, accompanied by the consolidation of digital communication ecosystems capable of sustaining mobilization and symbolic dispute during critical moments [Norris and Inglehart, 2019]. The far-right is a political ideology whose main ideals are nationalism, xenophobia, racism, misogyny, and anticommunism Rydgren [2007]. Far-right politics has diverse expressions, manifesting itself through openly neofascist organizations, others referred to as the alternative right, to more traditional groups integrated into institutional political engagement [Mudde, 2016; Pirro, 2022]. Its ascendancy is associated with profound social transformations, such as economic crisis, the weakening of liberal democratic institutions, and the emergence of new forms of social interaction, particularly digital platforms, social media networks, and instant messaging applications [Mudde, 2019].

In Brazil, WhatsApp is configured as the most widely used social network, being present in practically all smartphones in the country, according to data from Statista[1]. Its enormous popularity can be attributed to several factors, among which stand out the oral tradition of Brazilian culture, which favored the massive use of voice messages; the incentive from telecommunications operators, which exempt the application's data consumption from mobile data plans (a practice known as zero rating); and the user-friendly interface, which allows its use by people of practically all age groups, educational levels, and social classes. More than a messenger, WhatsApp should be viewed as a social network. Relations among users occur in chat groups and broadcast lists. Such groups are accessible via invitation links and have specific discussion themes, such as politics, sports, or finance. The main difference is that in groups, everyone can share messages, whereas in broadcast lists, messages are sent directly from administrators to members, without interaction between them. It allows a user to participate in hundreds of

[1] Available at: https://www.statista.com/topics/7731/whatsapp-in-brazil/#topicOverview. Accessed on: September 10, 2023

groups and interact with thousands of people. In 2023, a new function called "Community" was introduced in Brazil, allowing administrators to send information to multiple groups simultaneously. Through these features, a given piece of content can reach a large number of individuals in a short period of time [de Sá *et al*., 2023].

Given this context, the Brazilian far-right has begun to consolidate its own communication ecosystem, in which WhatsApp plays a prominent role. This structure has operated simultaneously on two fronts: on one hand, they promote attacks on democratic institutions; on the other, they seek to foster a political culture grounded in conservative values. Thus, to comprehensively understand the interactions and collective behaviors of far-right activists on WhatsApp, as well as how these actions shape information flow and opinion formation, it is fundamental to adopt analytical approaches that capture not only the content of messages but also the structure and dynamics of communication networks. To this end, complex network theory offers a robust methodological framework for analyzing interconnections among users and groups, identifying patterns in information dissemination, and detecting influential communities within the network [Evangelista and Bruno, 2019; Chagas and Da-Costa, 2023; Machado *et al*., 2019].

In our previous work [Nascimento *et al*., 2025], we conducted an exploratory analysis of the dynamics of WhatsApp groups associated with the Brazilian far-right in two distinct moments, more precisely the pre-electoral period (August 16, 2022 to October 30, 2022) and the post-electoral period (November 1, 2022 to February 15, 2023). Based on the collection of messages shared across 301 distinct public groups, bipartite networks were constructed between groups and messages, and their projections were subsequently used to construct group networks. In this manner, each investigated period was represented by a complex network. The analyses involved centrality measures, core-periphery decomposition, community detection, textual polarity analysis, and spatial visualizations. The results provided a macro view of the ecosystem's reorganization, i.e., how the "broad network" of groups reconfigures itself between pre- and post-electoral periods.

In this paper, we have extended our previous work, advancing beyond the macro view and focusing the analysis on synchronized activities to seek evidence of coordinated behavior. The concept of coordinated inauthentic behavior (CIB) constitutes a well-established operational paradigm within online influence campaigns, in which multiple accounts, frequently automated or inauthentic, amplify designated narratives through synchronized dissemination mechanisms ([Machado *et al*., 2019; Pacheco *et al*., 2021; Mannocci *et al*., 2024a; Tardelli *et al*., 2024]). Such accounts characteristically disseminate analogous content within compressed temporal intervals, repurpose media resources, or engage in reciprocal interactions to augment engagement metrics. Although coordination per se does not necessarily entail malicious intent, it becomes problematic when coupled with deceptive practices, specifically when inauthentic accounts masquerade as authentic individuals or exploit algorithmic mechanisms to artificially enhance visibility. These tactics are frequently deployed to advance particular political agendas, fabricate an appearance of consensus, or marginalize opposing perspectives, particularly during periods of heightened societal tension ([Vishnuprasad *et al*., 2024; Cinus *et al*., 2025]). The extant literature on CIB detection has predominantly concentrated on Twitter/X ([Pacheco *et al*., 2021]), where behavioral indicators such as co-retweets, shared uniform resource locators (URLs), and posting synchronicity have demonstrated considerable utility in identifying coordinated campaigns, including state-sponsored information operations. A comparatively limited corpus of scholarship has examined coordination mechanisms across alternative platforms, including Facebook and Telegram ([Cinus *et al*., 2025]). Conversely, the WhatsApp platform remains substantially underexamined, with scant prior investigations systematically interrogating CIB through computational methodologies on this particular medium. This gap assumes particular significance given WhatsApp's distinctive algorithmic architecture. In contrast to Twitter/X, where retweet networks and reply chains provide structural indicators of coordination, WhatsApp's design precludes the formation of explicit interaction graphs. Consequently, identifying coordination on WhatsApp requires methodological innovations calibrated to the platform's media format and engagement dynamics.

Thus, while our previous study describes the structural reorganization of the ecosystem as a whole, this work investigates "who coordinates with whom", with what intensity, and how these blocks transform from the pre- to the post-election period. As contributions, (i) we propose an operational pipeline to detect coordination on WhatsApp based on a combination of textual similarity and a temporal window; (ii) we model and analyze coordination networks among users in the pre- and post-election periods; and (iii) we contrast a broad layer of episodic coordination with a hard core of recurrent interactions through weight thresholds.

The remainder of this paper is organized as follows. Section 2 discusses the main related works. Section 3 describes the datasets developed and used. Section 4 describes the methods employed. In Section 5, the results obtained are presented and discussed. Finally, Section 6 presents the conclusions drawn and possibilities for future work.

2 Related Works

Contemporary scholarship in political communication increasingly acknowledges that structural network analyses alone are frequently insufficient to differentiate organic information diffusion from coordinated operations—specifically, instances in which multiple actors disseminate identical or near-identical content in synchronized temporal patterns to amplify message reach or artificially manufacture consensus. This recognition motivates methodological approaches that operationalize coordination through behavioral traces that jointly encode content repetition and temporal dynamics, rather than relying exclusively on aggregated co-occurrence patterns within extended temporal windows. Building on this theoretical framework, our investigation identifies coordinated events on WhatsApp by integrating short-term burstiness metrics with textual similarity analysis [Mannocci *et al*.,

2024b; Pacheco *et al*., 2021; Kansaon *et al*., 2025].

A prominent line of research models coordination by transforming behavioral traces into graph via behavioral traces that jointly encode repetition and time, rather than relying solely on aggregated co-occurrence across al. [Pacheco *et al*., 2021] propose a flexible, unsupervised framework for detecting coordination using short-term burstiness and coordination networks derived from diverse behavioral traces (e.g., temporal synchronization and content similarity), emphasizing that coordinated behavior may prove difficult to detect when accounts are examined in isolation and should therefore complement individual-level analytical approaches [Pacheco *et al*., 2021]. In earlier rapid-spread formulations, coordination is operationalized by enforcing constrained reaction times for diffusion actions (e.g., rapid retweets) and by constructing similarity networks derived from near-duplicate content.

In [Sharma *et al*., 2021], the authors proposed a technique to detect coordinated accounts based on their collective behaviour, inferred directly from their social media activities. The proposed method is independent of linguistic, metadata, or platform-specific features and can therefore be generalized across platforms, languages, and countries from which disinformation campaigns originate. The proposed approach was applied to the Russian Interference and COVID-19 datasets.

A pipeline-based approach to finding sets of accounts engaging in coordinated behaviour was proposed in [Weber and Neumann, 2021]. To reveal latent (i.e., hidden) networks of cooperating accounts, the authors proposed a temporal-window approach that relies solely on account interactions and metadata. It detects groups of accounts engaging in various behaviours that, in concert, execute different goal-based amplification strategies, a number of which we describe alongside other inauthentic strategies from the literature. The proposed approach relies upon a pipeline that extracts relevant elements from social media posts common to the major platforms, infers connections between accounts based on criteria matching the coordination strategies to build an undirected weighted network of accounts, which is then mined for communities exhibiting high levels of evidence of coordination using a novel community extraction method. The proposed approach was evaluated using two politically relevant Twitter datasets. Using discrete time windows, they temporally constrain potentially coordinated activities, successfully identifying groups operating over various time frames.

In [Magelinski *et al*., 2022], the authors proposed a multi-view, synchronized-action framework for analyzing coordination behavior across multiple actions. This general framework encompasses behaviors studied in previous works and allows for expansion into others. It is designed to uncover synchronized action within narrow time windows, allowing the detection of highly orchestrated covert campaigns as well as viral, public-facing campaigns, while avoiding investigation of coordination on the longer time scales characteristic of grassroots organizations. The authors demonstrated the effectiveness of the proposed approach using the Twitter dataset related to the Reopen America conversation. The proposed method demonstrated optimal scalability and processed three types of coordination across a nearly 10-million-node dataset in under 40 seconds.

In [Cinelli *et al*., 2022], the authors investigated the role of coordinated accounts in information spreading on Twitter, focusing on the 2019 UK political elections. In particular, they have considered a set of about 50,000 information cascades reconstructed from retweeting activity and the follower/following network. Overall, they found a small but significant difference in the behavior of coordinated and non-coordinated accounts in terms of node-level measures, namely: positioning level, action delay, and descendants count. Coordinated accounts tend to occupy higher positions in the information cascade, share information more rapidly, and infect more users. The results showed that the ability of coordinated accounts to involve noncoordinated ones follows a saturation-like process, characterized by a threshold beyond which injecting more coordinated accounts into a cascade yields negligible effects.

More recent work highlights that coordination is not necessarily a binary phenomenon. Nizzoli et al. introduce a non-binary, network-based framework that estimates coordination along a continuous range and iteratively uncovers coordinated communities across increasing coordination levels, avoiding a single fixed similarity threshold [Nizzoli *et al*., 2021]. This contrasts with threshold-based pipelines (including ours), where filtering is used to increase robustness and separate weak, episodic coordination from persistent patterns, a trade-off that is particularly relevant in high-volume settings.

Coordination detection has also been explored in the context of influence operations, astroturfing, and orchestrated campaigns. The survey presented in [Mannocci *et al*., 2024b] consolidates definitions and methodological families, providing taxonomies of coordination signals (text, URLs, timing, multimodality) and emphasizing open challenges related to temporal dynamics and cross-platform scarcity. Empirical studies further showed that coordinated behavior can manifest as synchronized amplification and propaganda diffusion through coordinated communities, reinforcing the importance of time-aware and content-aware signals [Hristakieva *et al*., 2022]. Additionally, the literature stresses that coordination and automation are orthogonal concepts, supporting the view that coordinated behavior should be studied independently of bot detection [Nizzoli *et al*., 2021; Keller *et al*., 2020].

In Tardelli *et al*. [2024], the authors carried out a dynamic analysis of coordinated behavior. They built a multiplex temporal network and applied a dynamic community-detection strategy to identify groups of users exhibiting coordinated temporal behavior. They found that: (i) coordinated communities feature variable degrees of temporal instability; (ii) dynamic analyses are needed to account for such instability, and results of static analyses can be unreliable and scarcely representative of unstable communities; (iii) some users exhibit distinct archetypal behaviors that have important practical implications; (iv) content and network characteristics contribute to explaining why users leave and join coordinated communities. Their results demonstrate the advantages of dynamic analyses and open new directions of research on the unfolding of online debates, the strategies of coordinated communities, and the patterns of online influence.

While most coordination detection research concentrates on Twitter/X due to data availability, fewer studies examine coordination on messaging platforms, where limited public visibility and collection constraints complicate analysis [Mannocci *et al*., 2024b]. In the Brazilian context, prior works have documented the centrality of WhatsApp to political mobilization and the migration of narratives from public arenas to opaque group-based environments, motivating scalable methods tailored to messaging logs [Machado *et al*., 2019; Evangelista and Bruno, 2019; Chagas, 2022]. Kansaon et al. provide empirical validation that coordinated campaigns on WhatsApp frequently operate within short burstiness windows, offering an important platform-aligned reference for time-aware coordination detection [Kansaon *et al*., 2025].

Table 1 summarizes representative coordination-detection approaches, highlighting differences in signals, graph construction, and temporal modeling, and clarifying how our RSN-based operationalization fills the WhatsApp-specific gap.

Our work is positioned at the intersection of (i) trace-to-graph coordination detection and (ii) messaging-platform political analysis. Different from macro-level studies that characterize group ecosystems or channel networks, we focus on *behavioral evidence* of synchronized repetition among users. Concretely, we adopt a Rapid Spread Network (RSN) operationalization: two messages are treated as a coordinated interaction when they are published by distinct users within a short time window and exhibit high textual similarity (TF–IDF + cosine), using empirically validated parameters (e.g., $\tau_{time} = 60$s, $\tau_{sim} = 0.70$). We then aggregate these events into a weighted user coordination network and contrast a broad layer of episodic coordination with a persistent hard core through weight thresholds, enabling a direct pre- vs. post-election comparison of coordination structure on WhatsApp [Pacheco *et al*., 2021; Kansaon *et al*., 2025; Mannocci *et al*., 2024a].

3 The Proposed Datasets

To analyze the coordinated behavior in public WhatsApp groups used by the Brazilian far-right during the 2022 pre- and post-election periods, we constructed two datasets, PRE2022 and POS2022. These datasets incorporate all messages circulated within the observed groups. In Table 2, we present basic statistics computed for the PRE2022 and POS2022 datasets, including the number of groups, users, messages, and unique messages.

Data collection was performed using the **BATMAN** (Big dAta Tools for Monitoring wAtsapp Networks) platform de Sá *et al*. [2023]. This tool enables automated monitoring of WhatsApp *public groups*. We selected public groups in which messages associated with the Brazilian far-right had circulated. It should be noted at the outset that not all users present in the observed public groups can be characterized as far-right extremists. Nevertheless, the ideas propagated by the Brazilian far-right (including racism, misogyny, xenophobia, LGBTphobia, authoritarianism, militarism, and religious fundamentalism, among others) were strongly preva-

Table 1. Main Features of Related Works.

Work	Platform/data	Signals	Time	Graph unit	Scope / output
Pacheco et al. (2021)	Social media (traces)	behavioral traces (e.g., time, content)	short-term burstiness / sync	coordination network	unsupervised framework
Sharma et al. (2021)	(2 datasets)	collective behavior; feature-agnostic	–	–	coordinated accounts
Weber & Neumann (2021)	Twitter	interactions + metadata	discrete windows	weighted account net	communities / pipeline
Magelinski et al. (2022)	Twitter	multi-view synchronized actions	narrow windows	multi-view action graphs	scalable detection
Cinelli et al. (2022)	Twitter (cascades)	cascade role metrics	cascade time / delays	cascades (+ network)	diffusion impact
Nizzoli et al. (2021)	(UK election context)	non-binary coordination	progressive levels	backbone + communities	continuum of coord.
Mannocci et al. (2024)	survey	taxonomy (text/URL/time/multi)	open challenges	–	survey / gaps
Tardelli et al. (2024)	temporal analysis	temporal coordination	dynamic modeling	multiplex temporal net	dynamic communities
Kansaon et al. (2025)	WhatsApp	burstiness evidence	short windows	network analysis	WhatsApp validation
This work (RSN)	**WhatsApp**	**text sim. + burstiness**	τ_{time}, τ_{sim}	**weighted user net**	**pre/post; hard core**

Table 2. Basic Statistics of the PRE2022 and POS2022 Datasets.

Statistic	Pre-election Dataset (PRE2022)	Post-election Dataset (POS2022)
Collection Period	August 16 to October 30, 2022	November 1, 2022, to February 15, 2023
Number of Groups	327	241
Number of Users	15,289	11,997
Number of Messages	260,153	276,748
Number of Unique Messages	165,520	186,149

lent in the content that circulated within the investigated groups. It is also worth emphasizing that, despite this work examining more than 320 public groups, it is not possible to obtain a comprehensive picture of the far right's use of WhatsApp. For instance, it was not feasible to determine precisely how many groups are, in fact, actively used by far-right actors in Brazil. Notwithstanding these limitations, it was possible to identify coordinated behavior in the Brazilian far-right's WhatsApp public groups.

The observed groups were identified through a systematic process consistent with the methodology established in the WhatsApp monitoring literature [Resende *et al*., 2019; Machado *et al*., 2019]: publicly shared invitation links, circulated by group administrators on social media platforms (Twitter/X, Facebook, Telegram) and indexed via web crawling by the BATMAN platform de Sá *et al*. [2023], were matched against a thematic dictionary comprising lexical markers associated with the Brazilian far-right ideological profile, specifically terms related to nationalism, militarism, anti-communism, and attacks on democratic institutions (particularly the judiciary and the electoral system). This profile follows the conceptual characterization of far-right ideology adopted in the political science literature [Rydgren, 2007; Mudde, 2019]. A group was included in the monitored corpus when its name, description, or publicly visible link matched one or more terms in the thematic dictionary, following an approach analogous to that of [Resende *et al*., 2019], in which group URLs were matched against a dictionary of politician names, party names, and words associated with political extremism. We acknowledge that this selection approach, based on publicly advertised groups and domain-informed thematic filtering, does not guarantee complete coverage of far-right WhatsApp activity, since private or unlisted groups remain inaccessible by design. This is a structural constraint shared by all prior work in this research line [Resende *et al*., 2019; Kansaon *et al*., 2025; Machado *et al*., 2019].

To protect user privacy and comply with the requirements of the General Personal Data Protection Law (LGPD), as well as WhatsApp's privacy policy[2], the collected data were anonymized by applying *hash* functions to identifying attributes, such as phone numbers and user and group names. Naturally, the anonymization of sensitive information imposes limitations on the semantic interpretation of the results. In particular, it becomes unfeasible to link specific narratives to identifiable users or to characterize communities discursively. Despite this restriction, the network structure and user interaction metrics retain sufficient information to analyze communication dynamics and structural changes across the analyzed periods.

4 Methodology

To investigate the coordinated behavior in WhatsApp groups, we developed an analytical *pipeline* composed of four sequential stages:

[2]https://www.whatsapp.com/legal/privacy-policy

- Text preprocessing and vectorization;
- Detection of coordinated behavior via sliding time windows;
- Complex network modeling and analysis;
- Semantic and content characterization.

The four stages described next are each designed to address a specific analytical layer of the research. Table 3 provides an overview of the techniques employed at each stage, their purpose within the broader argument, and the sections where results are reported. This mapping makes explicit how the detection pipeline (stage 2) feeds the structural analyses (stage 3) and how those structures are then interpreted through semantic characterization (stage 4).

Figure 1 illustrates the proposed pipeline. Next, we will describe in detail each one of the four stages of the proposed pipeline.

Besides, it is important to highlight that using the PRE2022 and POS2022 datasets described in Section 3 (Table 2), we applied the same methodological pipeline to both periods (2022 pre- and post-election). The following steps detail the preprocessing choices, the vectorization parameters, and the operational criteria used to identify coordinated interactions and construct the aggregated user network.

4.1 Preprocessing and Vector Representation

Given the informal nature of WhatsApp communication, we applied a cleaning step to the raw data from the PRE2022 and POS2022 datasets. The process consisted of:

1. **Normalization:** Converting all text to lowercase and standardizing characters to reduce superficial variations.
2. **Noise removal:** Removing excessive punctuation and non-alphanumeric special characters, along with filters to exclude empty or null messages after processing.
3. **Stopword removal:** Functional words were removed using the extended NLTK library list for Portuguese, supplemented by recurrent terms that lacked semantic utility for the analysis.
4. **Entity preservation and normalization:** We opted to preserve informative content. URLs and numbers were textually normalized (e.g., removing `https://` protocols) but maintained within the message so that links and specific values would contribute to the similarity calculation. Emojis were converted to their textual representations (demojized), enabling visual symbols to be processed by similarity algorithms.

Vector Representation (TF–IDF). The preprocessed messages were represented by TF–IDF vectors using `TfidfVectorizer(min_df=2, dtype=np.float32)` globally for each period, generating a sparse matrix in which each message m_i is represented by a vector v_i. Since TF–IDF applies L_2 normalization by default, the similarity between two messages was estimated via cosine similarity, computed efficiently as the dot product between sparse vectors:

$$\mathrm{Sim}(v_i, v_j) = \frac{v_i \cdot v_j}{\|v_i\|\|v_j\|}.$$

4.2 Detection of Coordinated Behavior

Coordination detection is based on the premise of "surprising lack of independence" [Mannocci *et al*., 2024b]. We define coordination as the occurrence of messages with identical or highly similar content disseminated by distinct users within a short time interval.

It is important to note that this operational definition functions at the *behavioral level*: it identifies statistically improbable patterns of textual co-occurrence among distinct users, consistent with what the literature characterizes as a "surprising lack of independence" [Mannocci *et al*., 2024b; Nizzoli *et al*., 2021]. As recognized across the coordination detection literature, computational pipelines based on content similarity and temporal proximity cannot, by themselves, distinguish between *conjoint action*, in which behaviorally similar outputs emerge independently from shared ideological motivation, and *coordinated action* in the stronger sense, which presupposes an explicit relational structure, shared protocol, or command among agents [Rogers and Righetti, 2025; Giglietto *et al*., 2020]. This distinction is an acknowledged epistemological boundary of the entire methodological family [Cinelli *et al*., 2022; Thiele *et al*., 2025]. Accordingly, the structural detection stage presented here should be understood as a *first-stage filter* that surfaces behavioral patterns warranting further investigation. The semantic and content characterization stage (Section 4.4) serves as the interpretive second layer, contextualizing detected patterns within the political and informational dynamics of the analyzed ecosystem and supporting more grounded inferences about the nature of the observed coordination.

Figures 2a and 2b illustrate the cumulative distribution of textual similarity between message pairs in the PRE2022 and POS2022 periods, considering a 60-second time window. In both periods, most interactions are concentrated at low similarity values, while a small fraction occupies the high-similarity tail. This pattern supports selecting an intermediate similarity threshold that captures subtle content differences without requiring exact copies. In general, an inverse relationship exists between similarity and message quantity. Specifically, there is a considerable prevalence of message pairs exhibiting low similarity. Nevertheless, it is important to emphasize that the number of message pairs classified as similar increases sharply beyond the 0.7 similarity threshold.

The choice of $\tau_{time} = 60$s and $\tau_{sim} = 0.70$ was guided by the empirical distributions shown in Figure 2, which together constitute the parameter validation analysis. For τ_{time}, we adopt the *sequential distance* offset k as a robustness probe: for a chronologically ordered message stream, k indicates that m_j is the k-th message following m_i in the global stream (e.g., $k = +1$ for the immediately next message, $k = +2$ for the second next, up to $k = +10$). The left-hand panels of Figure 2 show the complementary cumulative distribution function (CCDF) of time intervals Δt for $k \in \{+1, +2, +3, +4, +5, +10\}$: all curves exhibit a sharp drop within the first tens of seconds and converge to a similar shape across offsets, indicating that the burstiness pattern is not a methodological artifact of adjacent-message comparisons but a stable, stream-wide property of the data. Values of τ_{time} between 30s and 120s were also evaluated; in all cases

Table 3. Analytical Pipeline: Stages, Techniques, and Purposes.

Stage	Technique	Purpose in the argument
(1) Preprocessing & Vectorisation	Text normalisation; stopword removal; TF–IDF	Reduce noise while retaining semantically informative tokens; enable efficient cosine similarity computation
(2) Coordination Detection	Cosine similarity + sliding window (τ_{sim}, τ_{time})	Identify message pairs with statistically improbable co-occurrence — the behavioural coordination signal
	Weight filter ($w \geq 2$; $w \geq 5, 10, 20$)	Separate episodic coincidence from persistent, recurrent coordination
(3) Network Modelling & Analysis	Weighted graph construction	Aggregate pairwise coordination events into a user-level network for pre- vs. post-election comparison
	K-shell (k-core) decomposition	Identify the *hard core* of recurrent coordinators; quantify structural intensification across periods
	Louvain community detection	Uncover latent coordination "teams"; assess their organisational stability and modular structure
	Centrality (strength, betweenness, eigenvector)	Characterise division of labour between internal amplifiers and inter-cluster bridges
(4) Semantic & Content Characterisation	BERTopic per community	Contextualise structural detection; identify dominant narratives within each coordination team
	Sentiment + misinformation score	Distinguish incidental coordination from politically oriented content
	Geospatial inference (area codes)	Map approximate geographic distribution of the coordination core across periods

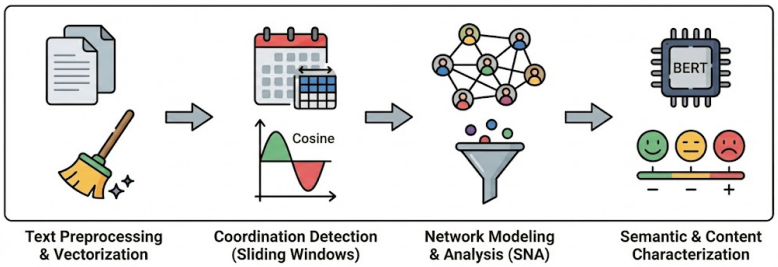

Figure 1. Methodological Pipeline.

the sharp drop remained well within the chosen 60-second window, confirming that the qualitative conclusions and the pre- vs. post-election comparison are robust to moderate variations in the time threshold. The value of 60s was retained as it is consistent with the burstiness window empirically validated for coordinated campaigns on WhatsApp by Kansaon *et al*. [2025].

For τ_{sim}, the right-hand panels of Figure 2 show the probability density of textual similarity for pairs published within $\Delta t < 60$s. The distribution presents a pronounced density valley in the interval $[0.7, 0.85]$, followed by a sharp peak near similarity $= 1$, corresponding to near-exact copy–paste behavior. The threshold $\tau_{sim} = 0.70$ was placed at the onset of this valley, capturing both exact copies and messages with slight variations (e.g., minor reformulations or obfuscation attempts) while excluding the large mass of weakly similar pairs concentrated at low similarity values. Setting τ_{sim} at a higher value (e.g., 0.85) would restrict detection to near-exact copies only, missing the intentional variation typical of coordinated obfuscation; a lower value (e.g., 0.60) would risk incorporating thematically similar but independently authored messages, increasing false positives. The chosen threshold, therefore, represents an empirically grounded balance between sensitivity and specificity, supported by the bimodal shape of the similarity distribution.

In the left-hand panels of Figure 2, we show the CCDF of the time intervals Δt between message pairs with textual similarity Sim ≥ 0.7, considering different sequential distances $(+1, +2, +3, +4, +5, +10)$. A sharp drop in the CCDF is observed within the first tens to hundreds of seconds, indicating that a significant fraction of highly similar messages is disseminated in very short intervals. The similarity in the shape of the curves for different offsets suggests temporal stability and the recurrence of the synchronized propagation process.

The right-hand panels of Figure 2 present the textual similarity distribution (PDF) for message pairs published within $\Delta t < 60$ seconds. A pronounced density peak is observed for similarity values near 1, indicating near-exact textual replication typical of copy–paste behaviors or re-posts with minimal modifications. At the same time, the existence of probability mass at values slightly below 1 justifies the adoption of $\tau_{\text{sim}} = 0.7$, allowing for the capture of slight variations intentionally introduced for obfuscation without compromising the specificity of the criterion.

Taken together, the patterns observed in the temporal (CCDF) and semantic (PDF) distributions for PRE2022 and POS2022 provide empirical support for the operational definition of coordination adopted in this work, as well as for using a short time window to distinguish coordinated events from organic interactions distributed over time.

Following the *framework* proposed by Pacheco *et al*.

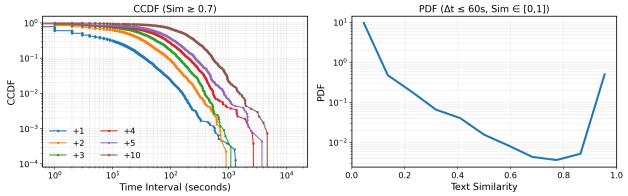

(a) PRE2022

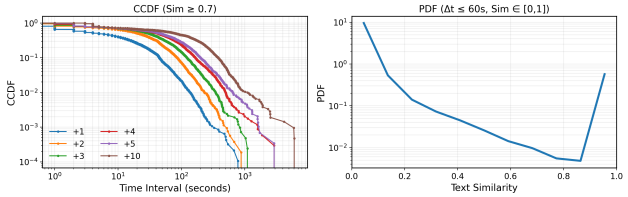

(b) POS2022

Figure 2. Temporal and Semantic Analysis in the Rapid-spread Style. **Left**: CCDF of time intervals Δt between message pairs with textual similarity ≥ 0.7, computed for sequential distances $k \in \{+1, +2, +3, +4, +5, +10\}$, where k denotes the positional offset between two messages in the chronologically ordered global stream (e.g., $k = +1$ compares each message with the immediately following one). The convergence of curves across offsets indicates that the observed burstiness is a robust, stream-wide pattern rather than an artefact of adjacent-message comparisons. **Right:** PDF of textual similarity for pairs published within $\Delta t < 60$ seconds.

[2021] and empirical validations reported by Kansaon *et al.* [2025], we adopted a *Rapid Spread Network* (RSN) approach to detect coordinated behavior. Let m_i and m_j be two messages sent by distinct users $u_a \neq u_b$ at timestamps t_i and t_j, represented by TF–IDF vectors v_i and v_j. We define $\Delta t = |t_i - t_j|$. A message pair is considered a **coordinated interaction** if it simultaneously satisfies:

$$\text{Sim}(v_i, v_j) \geq \tau_{sim} \quad \text{e} \quad \Delta t \leq \tau_{time}. \tag{1}$$

- $\text{Sim}(v_i, v_j)$ is the cosine similarity between the message vectors. We define $\tau_{\text{sim}} = 0.70$, allowing for the detection of not only exact copies but also slight variations.
- τ_{time} is the sliding time window applied to the global stream of the period. We set $\tau_{\text{time}} = 60$ seconds. This choice is consistent with evidence that coordinated campaigns on WhatsApp operate predominantly within this 1-minute burstiness window [Kansaon *et al.*, 2025], in contrast to organic diffusion on scales of hours or days.

Temporal Indexing and Computational Cost. To avoid quadratic comparisons among all pairs of messages, the messages were chronologically ordered. For each message m_i with timestamp t_i, we use binary search (*bisect*) to locate the candidate interval $[t_i, t_i + \tau_{\text{time}}]$ (sliding window) and compute textual similarity only within this subset. This strategy reduces the computational cost from $O(n^2)$ to approximately $O(n \cdot k)$, where k represents the average number of messages within the time window.

4.3 Complex Network Modeling and Analysis

User-aggregated network. We constructed an undirected and weighted graph $G = (V, E, W)$ for each period, where each node represents a user and each edge represents coordination between a pair of users, regardless of the messages' source group. Each detected coordinated interaction between users (u_a, u_b) increments the weight w_{ab} by one unit, so that w_{ab} reflects the total number of coordinated messages between u_a and u_b throughout the period.

To increase the robustness of the analysis and reduce random coincidences, we applied a minimum-weight filter: only edges with $w_{ab} \geq 2$ were retained. Additionally, we performed robustness analyses considering more restrictive weight thresholds (e.g., $w \geq 5, 10, 20$) to contrast the broad network "with the coordination hard core".

Based on the filtered graph, we calculated metrics to characterize the coordinated behavior:

- **Coordination weight (W_{coord}):** the sum of the weights of the edges incident to a node (strength), representing its total volume of synchronized activity.
- **Capillarity (reach across groups):** the number of distinct groups (group IDs) associated with coordinated message occurrences, allowing for the distinction between more local actors and those with broader reach across groups.
- **Core–periphery decomposition:** In k-core decomposition, each node receives a shell index equal to the maximum k for which it belongs to the k-core subgraph; nodes in the outermost shell (index k_{max}) form the strict core, while nodes in the immediately preceding shell (index $k_{max} - 1$) constitute the second layer. We define the *global core* as the union of these two layers — i.e., users with coreness $\geq k_{max} - 1$ (top-2 shells) — while all lower shells represent the periphery. .

To identify groups acting in concert (coordination "teams"), we applied the Louvain algorithm for community detection (modularity maximization) on the weighted graph, utilizing the `Weight` edge attribute [Zheng *et al.*, 2024]. The node–community partitioning was then used as a basis for semantic analyses and structural comparisons between periods.

Additionally, we calculated centrality metrics to identify specific roles within the network [Iswarya *et al.*, 2024]:

- **Strength / Weighted Degree:** measures the total intensity of a user's coordinated activity.
- **Degree (unweighted):** measures the number of distinct neighbors in the filtered graph.
- **Betweenness Centrality:** identifies users who act as bridges between communities; it was computed in normalized form.
- **Eigenvector Centrality:** identifies influential users connected to other highly active individuals; it was computed considering edge weights.

4.4 Semantic and Content Characterization

To understand the content disseminated by the detected communities, we utilized Natural Language Processing (NLP) techniques on the preprocessed text.

Topic Modeling (BERTopic). The thematic characterization was performed using **BERTopic** on the corpus of messages from users present in the network (obtained via an inner

join between the message dataset and the community file). Messages with null or empty text after preprocessing were discarded.

The topic distribution per community was obtained via `topics_per_class` (class = `community_id`). We used dimensionality reduction via UMAP with `n_neighbors=15`, `n_components=5`, `min_dist=0.0`, `metric=cosine`, and `random_state=42`, followed by clustering with HDBSCAN (`min_cluster_size=15`, `metric=euclidean`, `cluster_selection_method=eom`, `prediction_data=True`). The model was configured with `language="multilingual"` and `min_topic_size=10`.

Topic representation post-processing. After training, we updated the textual representation of the topics to reduce tokenization noise using `update_topics`, defining a `CountVectorizer` with `token_pattern=\b\w{2,20}\b`. This adjustment restricts the terms considered in the c-TF–IDF composition to alphanumeric tokens between 2 and 20 characters, increasing the interpretability of the keywords associated with each topic. The BERTopic/HDBSCAN outlier topic (`Topic=-1`) was excluded from the interpretive analyses. The resulting topics were interpreted based on the most representative words (c-TF–IDF) and message examples, allowing for the description of dominant narratives within each community (e.g., "voter fraud", "military intervention", "general strike").

Sentiment and Misinformation Analysis. Sentiment polarity was estimated using *TextBlob* and discretized into positive (polarity > 0.1), negative (polarity < -0.1), and neutral (otherwise). To assess misinformation, we used the `score_misinformation` metadata provided by the BATMAN platform[3]. and aggregated statistics by community (e.g., the mean), allowing us to distinguish clusters with potentially benign coordination (e.g., spam) from clusters associated with political misinformation campaigns.

Geospatial Inference. We inferred the approximate location from metadata (area codes [DDD] and/or latitude/longitude) aggregated by user. For each user, we utilized the median latitude/longitude and the mode of the area code (DDD)/state/country across their messages, cross-referencing these variables with the core/periphery and the communities for a spatial analysis of the command structure between the periods, using the `Folium`[4].

All analyses were implemented using the *NetworkX*[5], *PyVis*[6], *Folium*, and *TextBlob* libraries, as well as the *BATMAN* platform. The code and the results are available in our *online* repository[7].

[3] `https://faroldigital.info/`

[4] `https://python-visualization.github.io/folium/`

[5] `https://networkx.org/`

[6] `https://pyvis.readthedocs.io/`

[7] https://github.com/jmmfilho/complexnetworks-whatsapp-2022elections

5 Experimental Results

5.1 Detection of Coordinated Behavior

Based on the coordinated behavior detection threshold ($\tau_{\text{sim}} = 0.7$, $\tau_{\text{time}} = 60$s) and the robustness filter ($w \geq 2$), we constructed two weighted coordination networks: one for PRE2022 and another for POS2022. Table 4 summarizes the structural properties of these networks. In Table 4, column $|V|$ indicates the number of vertices in the network, while column $|E|$ represents the number of edges. Besides, column $\sum w$ displays the sum of the weights of all edges in the network and column $\bar{w}$ shows the average weight per edge ($|\sum w|/|E|$). Furthermore, column %(w = 2) indicates the percentage of edges whose weight is exactly 2 (the smallest possible weight, that is, a lower bound with respect to weight). The column %($w \geq 5$) displays the percentage of edges with weight greater than or equal to 5 (that is, edges that indicate a strong association between vertices). Finally, column *Dens* represents the density of the graph, that is, the fraction between the number of existing connections and the maximum possible number of edges ($2|E|/|V|(|V| - 1)$). Values close to zero indicate a sparse network, that is, one with low node association. The topological visualizations of these structures are presented in Figure 6, illustrating the connectivity and the distribution of clusters in both periods.

The results show a significant reconfiguration in the post-election period. Although both networks remain sparse and fragmented, as visually confirmed by the multiple small components in Figure 6, POS2022 exhibits a larger size ($|V|$ and $|E|$) and, above all, a higher coordination intensity ($\bar{w}$). As shown in Table 4, $\sum w$ and $\bar{w}$ increase substantially, as does the fraction of strong connections ($w \geq 5$). Operationally, this indicates that post-event coordination occurs more frequently between user pairs, reinforcing the hypothesis of synchronized routines sustained by a more active core, as represented by the denser clusters in Figure 6b.

5.2 Complex Network Modeling and Analysis

5.2.1 Core–Periphery Decomposition

To quantify the structural hierarchy of coordination, we applied k-shell decomposition on the coordination networks. We define the *strict core* as the set of users in the maximum k-shell ($k = k_{\max}$). This definition captures the "hard core" of coordination—that is, the actors who are most "deeply" embedded in the connectivity structure.

Table 6 summarizes the structural properties of the networks resulting from the k-shell decomposition. Column *kmax* indicates the maximum number of cores obtained after the decomposition process. Column (*k=kmax*) shows the number of users in the innermost core. Additionally, column %Core displays the percentage of users in the innermost core. Finally, column "Persistent in Core" indicates the persistence of the core across the two analyzed periods, that is, how many users remained in the core in both networks.

Table 6 highlights a striking contrast between the periods: despite $k_{\max}$ being lower in POS2022, the proportion of users in the strict core increases significantly (from 2.35% to 10.97%). Furthermore, within the subset of users present in

Table 4. Structural Summary of the Coordination Networks.

Period	$\|V\|$	$\|E\|$	$\sum w$	$\bar{w}$	$\%(w = 2)$	$\%(w \geq 5)$	Dens.	Comp.	Largest comp.
PRE2022	170	129	721	5.59	67.4%	13.2%	0.0090	47	46
POS2022	237	176	2080	11.82	57.4%	23.3%	0.0063	71	59

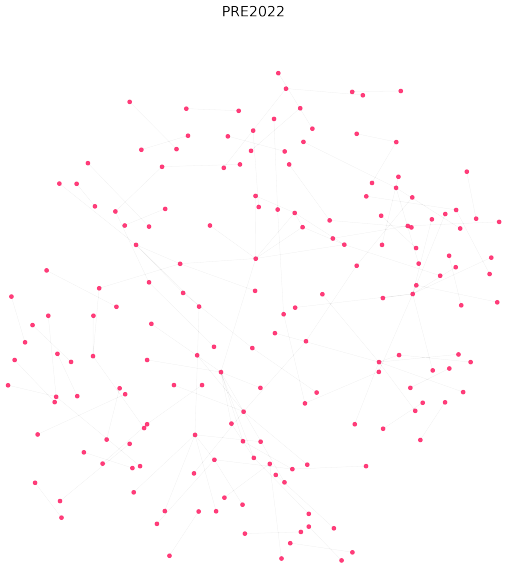

(a) PRE2022

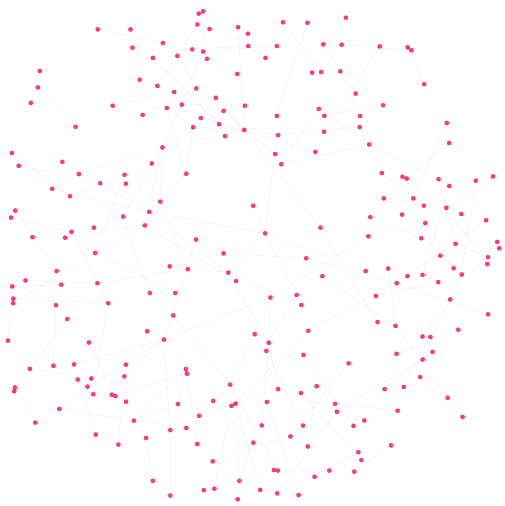

(b) POS2022

Figure 3. Coordination Networks. Visualization of the coordination networks for the (a) PRE2022 and (b) POS2022 periods. Nodes represent unique users and edges indicate detected coordinated interactions.

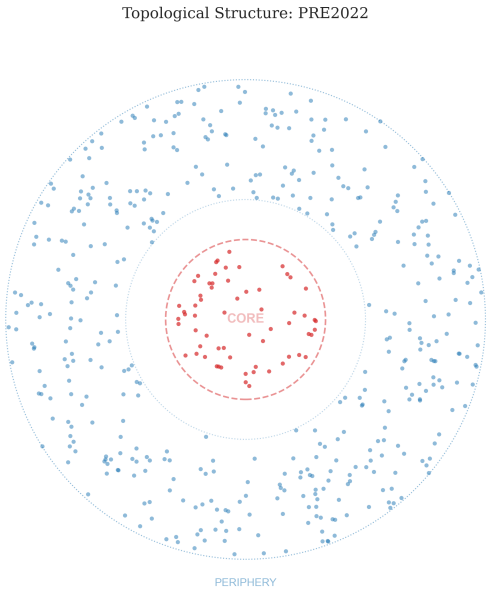

(a) PRE2022

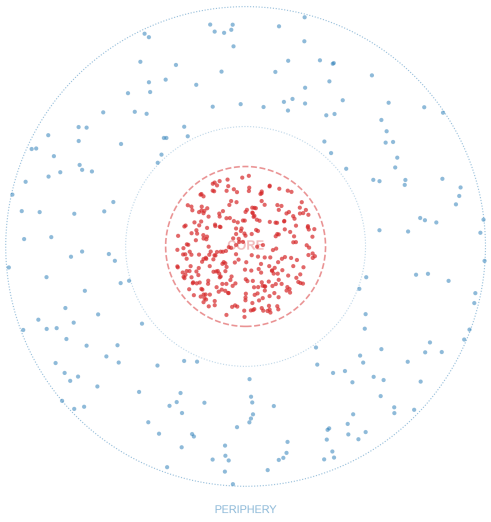

(b) POS2022

Figure 4. Core–Periphery Decomposition of Coordination Networks. Visualization of the k-shell decomposition applied to the coordination networks for the (a) PRE2022 and (b) POS2022 periods. Nodes represent unique users and edges indicate detected coordinated interactions.

both periods, the fraction occupying the strict core is higher in POS2022 (7/36) than in PRE2022 (3/36), indicating that a larger share of the comparable, persistently active participants integrates the coordination hard core in the post-event network. This pattern is consistent with a process of structural intensification in the post-election period, where more participants become part of the coordination core rather than remaining in peripheral layers.

5.2.2 Community Detection

To understand the structural organization of users in the two analyzed periods, the Louvain community detection algorithm was applied to the projected coordination networks. Communities can be understood as groups of nodes that are more densely connected to each other than to the rest of the network. Communities represent latent structures or natural groupings within the network, revealing patterns of organiza-

tion and behavior. Figure 5 illustrates the projected networks of the two periods, with distinct coloring for each community identified by the Louvain algorithm, highlighting how the users anchor the main coordination "teams" during the election cycle.

Table 5 summarises the structural properties of the community partition obtained via the Louvain algorithm on the weighted coordination networks ($w \geq 2$) for both periods. The absence of singletons in both periods is a direct consequence of the minimum-weight filter: every node retained in the network holds at least one coordination edge, so the algorithm assigns all nodes to communities of size $|V| \geq 2$.

Table 5. Community Detection Summary (Louvain, $w \geq 2$). Top-3 community sizes: PRE2022 = {13, 13, 12}; POS2022 = {30, 29, 7}.

Period	#Comm.	Q	Largest $\|V\|$	Pairs
PRE2022	51	**0.779**	13	34
POS2022	72	**0.745**	30	52

Both periods exhibit strongly modular organization. Modularity values of $Q = 0.779$ (PRE2022) and $Q = 0.745$ (POS2022) are well above the conventional significance threshold of $Q > 0.3$ and approach the empirical upper bound ($Q \approx 0.8$) Newman [2006], indicating that coordination "teams" are clearly delineated from one another in both networks. The slight decrease in modularity in the post-electoral period is consistent with the broader increase in coordination intensity documented in Table 4: as edge weights grow and more users enter the coordination structure, a modest number of inter-community edges emerges, marginally reducing compartmentalization without dissolving the overall community architecture.

The size distribution of communities is markedly right-skewed in both periods, dominated by coordination pairs. In PRE2022, 34 of 51 communities (66.7%) have exactly 2 users, while the remaining 17 range from 3 to 13 members. In POS2022, the proportion of pairs rises further: 52 of 72 communities (72.2%) are dyads, yet the non-dyadic communities become substantially larger, with the three biggest reaching 30, 29, and 7 users respectively. This pattern, a proliferation of small peripheral dyads alongside the consolidation of a few large, densely connected teams, mirrors and complements the core–periphery structure described in Section 5.2.1, where a growing fraction of users migrates into the innermost k-shell after the election. Together, these findings suggest a dual dynamic in the post-electoral period: an expansion of the coordination ecosystem at its periphery and a simultaneous concentration of activity within a reduced set of dominant communities.

Figure 5 illustrates the community partition for both periods; node size is proportional to coordination weight (W_{coord}) and node color encodes community membership. Given the sparsity of the networks, small communities and isolated pairs may be difficult to distinguish in the static layout. An interactive version of these visualizations, with adjustable layouts, zoom, and per-node metric display, is available via the project dashboard.[8]

[8] https://fernandanascimento26.github.io/Visualizacao-de-dados/

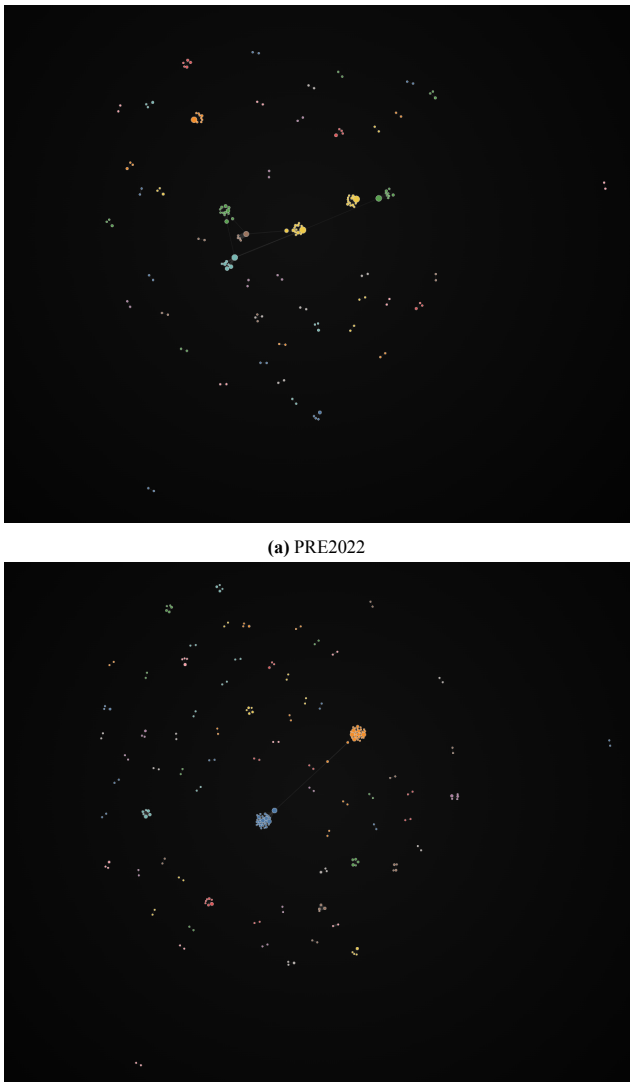

(a) PRE2022

(b) POS2022

Figure 5. Detected Communities in the Coordination Networks. Visualization of detected communities in the coordination networks for (a) PRE2022 and (b) POS2022. Node colors represent distinct communities identified via the Louvain algorithm.

5.2.3 Central Actors and Structural Roles

To identify relevant actors and differentiate their roles within the coordination, we analyzed complementary metrics: (i) strength (weighted degree), as a proxy for the total volume of synchronized activity; (ii) betweenness centrality, to detect bridge users; and (iii) eigenvector centrality, as a measure of influence surrounding the core.

Before inspecting the obtained rankings, we evaluated the relationship between these measures. Table 7 shows a very high association between strength and eigenvector centrality in both periods, indicating that actors with higher coordinated activity also tend to be the most influential within the structure. On the other hand, the correlation between eigenvector and betweenness centrality is close to zero, suggesting a functional separation: the influential core generally does not coincide with the primary connectors between clusters. This dissociation is significant because it points to an organization typical of coordinated systems: a set of actors highly connected to each other and a distinct subset responsible for disseminating information or linking different clusters.

Table 6. Core–Periphery Decomposition Summary.

Period	$k_{\max}$	Core ($k = k_{\max}$)	% Core	Persistent in Core
PRE2022	3	4/170	2.35%	3/36 (8.33%)
POS2022	2	26/237	10.97%	7/36 (19.44%)

Table 7. Correlation (Pearson r) between Centrality Measures in the Coordination Networks ($w \geq 2$).

Metric Pair	**PRE2022**	**POS2022**
Strength (Weighted Degree) × Influence (Eigenvector)	0.905	0.888
Degree × Betweenness	0.748	0.678
Influence (Eigenvector) × Betweenness	0.125	-0.027
Degree × Influence (Eigenvector)	0.170	0.196

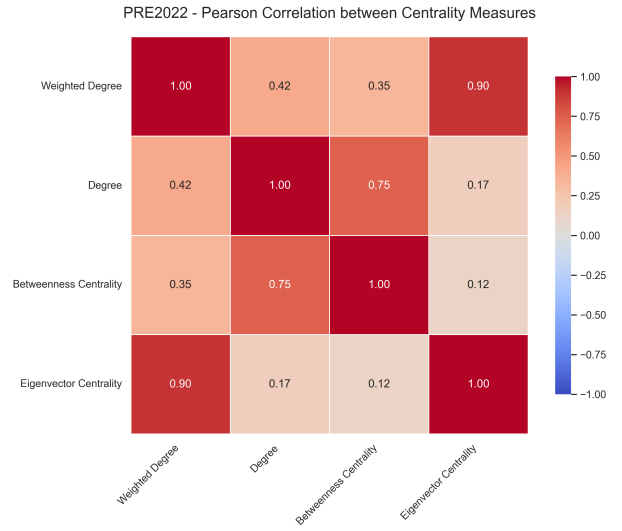

(a) PRE2022

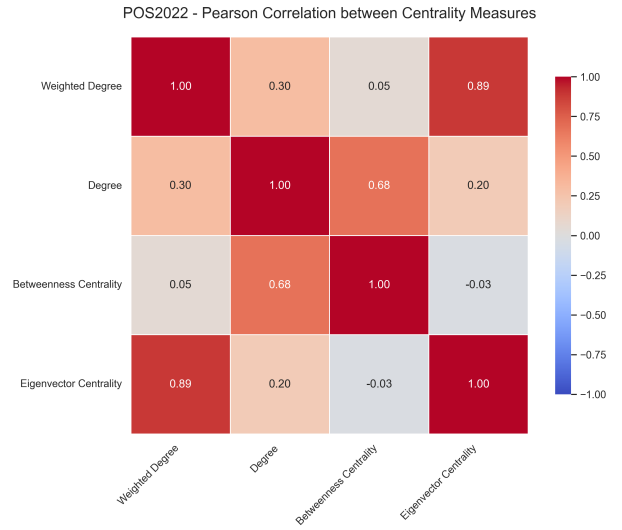

(b) POS2022

Figure 6. Person Correlation between Centrality Measures.

The centrality rankings corroborate this interpretation. In PRE2022, few users simultaneously dominate both strength and eigenvector centrality, characterizing a more concentrated core. In POS2022, strength reaches even higher levels, while eigenvector influence remains concentrated, suggesting a core with a high capacity for internal amplification. Concurrently, the top nodes by betweenness tend to be distributed across specific communities, consistent with the role of connecting clusters and increasing transversal reach.

Table 8. Top-5 users by strength (weighted degree), interpreted as the total volume of coordinated activity (W_{coord}). Truncated IDs preserve anonymity.

PRE2022		
User	**Comm.**	**Value**
4ea1b3ee	6	260
ef496106	6	247
cebd2107	2	127
b6242381	2	72
19654cb7	2	60
POS2022		
User	**Comm.**	**Value**
ef496106	2	764
4ea1b3ee	2	652
fac2d68f	1	342
b66e161a	2	261
f9a3713e	1	254

Table 9. Top-5 Users by Betweenness Centrality (Bridge Actors). Truncated IDs preserve anonymity.

PRE2022		
User	**Comm.**	**Value**
36de5147	23	0.0409
cebd2107	2	0.0321
32bc3dcc	42	0.0269
c5fff63b	20	0.0226
cd45e9e7	42	0.0197
POS2022		
User	**Comm.**	**Value**
b6242381	13	0.0396
36de5147	39	0.0322
726b2c36	39	0.0303
29ce0928	39	0.0303
c5fff63b	39	0.0280

5.3 Semantic and Content Characterization

5.3.1 Dominant Narratives by Community

Following the structural detection, we investigated the content disseminated within each community (Louvain) by applying BERTopic per cluster. Tables 11 and 12 summarize

Table 10. Top-5 Users by Eigenvector Centrality (Core Influence). Truncated IDs preserve anonymity.

PRE2022		
User	**Comm.**	**Value**
4ea1b3ee	6	0.7071
ef496106	6	0.7069
218131c9	23	0.0114
384dbf12	6	0.0086
e5cbbbb4	42	0.0057
POS2022		
User	**Comm.**	**Value**
ef496106	2	0.6787
4ea1b3ee	2	0.6538
b66e161a	2	0.2991
3bc96453	2	0.1501
eb57badf	2	0.0023

the dominant topics in the largest communities of PRE2022 and POS2022 (top-3 by $|V|$). In the first column, the notation `ID` ($|V|$) indicates the community identifier and, in parentheses, its size (number of users/nodes).

In PRE2022, the dominant topics reflect patterns consistent with agenda-setting and grassroots mobilization during the campaign, featuring references to specific actors/institutions and the circulation of attack/justification content. In POS2022, a dominant community stands out for its heavy concentration on topics related to the circulation of links/articles and calls to forward content (an amplification pattern), alongside topics related to post-election reactions and surveillance/contestation. Importantly, structural detection is topic-agnostic; therefore, the characterization stage (topics, misinformation) is essential for distinguishing incidental coordination from coordination involving politically oriented content.

5.3.2 Super-dissemination and Capillarity

To distinguish between coordination intensity and multi-group reach, we analyzed super-spreaders by combining: W_{coord} (total coordinated activity), the number of repetitions of a prototype message, and capillarity (the number of groups reached by the user, normalized by the total number of groups in which their community was active). Table 14 summarizes the primary cases.

The results suggest two complementary profiles: (i) actors with very intense coordination concentrated in a few groups, consistent with a focused operation; and (ii) actors with moderate intensity but high capillarity, consistent with an inter-group distribution/amplification function. This contrast helps to interpret coordination as a multi-level phenomenon: a core that produces/repeats content and a functional periphery that amplifies its reach.

5.3.3 Misinformation by Community

To evaluate the association between coordination and informational quality, we aggregated the disinformation score by community. In PRE2022, the volume-weighted average of scored messages was 0.349 (total of 8,659 scored messages); in POS2022, it was 0.313 (total of 11,220). Despite the decrease in the weighted average, distributions by community remain heterogeneous, as illustrated in Figure 7, with clusters presenting elevated averages and a sufficient volume of scored messages. This reinforces the importance of community-level analyses rather than relying solely on global aggregates, enabling the identification of radicalization niches that would go unnoticed in the general average.

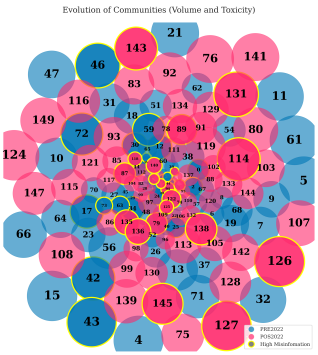

Figure 7. Heterogeneity and Volume by Community. Circle Packing visualization of the detected clusters. Bubble size is proportional to message volume, and color indicates the period: **Blue (Pre-Election)** and **Red (Post-Election)**. Bubbles with a **yellow border** highlight communities with high toxicity (Score ≥ 0.7), demonstrating the existence of radicalized echo chambers independently of the global average.

5.3.4 Geographical Distribution of the Networks' Core

Figure 8 illustrates the distinct geographical patterns between the two periods. In the Pre-Election period (left, blue scale), the coordination core is geographically diffuse, with São Paulo (SP) acting as the primary hub (9 users), followed by a dispersed presence in the Northeast (PB, RN) and Southeast (RJ). This suggests a multi-regional operation focused on broad dissemination during the campaign.

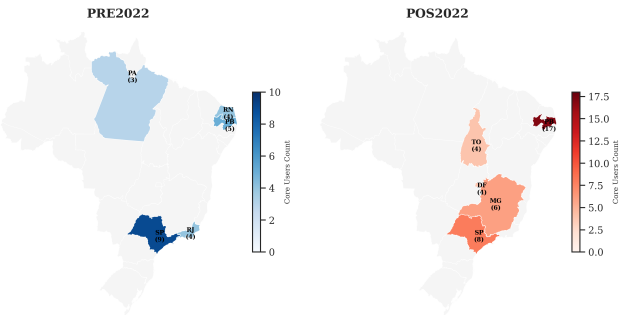

Figure 8. Geographical Distribution of the Coordination Core. Inferred from area codes (DDD). The left map (**Pre-Election**) shows a diffuse core with São Paulo as the main hub. The right map (**Post-Election**) shows a high concentration of core actors in Paraíba (PB), indicating a strategic geographical shift in coordination efforts.

In sharp contrast, the Post-Election period (right, red scale) reveals a significant shift towards high concentration in a single state: Paraiba (PB), which aggregates 17 users in the core, more than double the count of the second-placed state (SP, with 8). This indicates a strategic redirection of coordination efforts to a specific hub after the election results, possibly to align narratives or mobilize local actions more effectively.

Table 11. Dominant Topics (BERTopic) in the Largest Coordination Communities in PRE2022. The % column indicates the topic's share of the community's message volume.

Community ($\|V\|$)	Topic	Keywords (Top-5)	% in community
23 (13)	14	dallagnol, amplas, tre, investigação, deltan	2.70
23 (13)	34	remova, inédito, debate, foge, liga	2.27
23 (13)	32	agenda, destefani, semanal, 2310, participará	2.23
28 (13)	5	rádios, inserções, rádio, funcionário, 154	2.74
28 (13)	0	mula, justificar, cometendo, ptlixo, abs	2.60
28 (13)	27	números, 144, praticamente, cresce, pontos	2.46
9 (12)	13	canais, show, inscritos, pânico, curta	3.47
9 (12)	10	jovens, menor, general, prova, verdade	3.28
9 (12)	0	youtube, vídeo, assistam, canal, ao vivo	3.24

Table 12. Dominant Topics (BERTopic) in the Largest Coordination Communities in POS2022. The % column indicates the topic's share of the community's message volume.

Community ($\|V\|$)	Topic	Keywords (Top-5)	% in community
39 (30)	2	matérias, ler, links, acessar, conseguindo	8.98
39 (30)	1	matéria, pushpin, aliados, clicar, adicione	8.12
39 (30)	0	47992677688, adicione, conseguindo, policecarl..., ...	7.31
13 (29)	13	ministros, afetam, mudam, esquece, vigiar	1.61
13 (29)	22	indicam, 2024, depredações, duríssimo, exaltar	1.25
13 (29)	0	duck, repasse, esquerdopatas, pôr, adiante	1.16
41 (7)	108	folhadestracom, negou, portavozes, interditam,...	4.86
41 (7)	0	ouçam, seara, parapensarerefletir, tias, puder...	4.26
41 (7)	182	luiz, gauíba, beck, abreu, toni	3.04

The marked difference in concentration patterns between the periods provides strong evidence of a structural reorganization of the coordination network.

Table 15 summarizes the most frequent states and area codes within the expanded core, suggesting that coordination involves actors distributed across multiple regions, with shifts in relative concentration between periods. We emphasize that the area code is an imperfect proxy and may reflect limitations in registration or extraction; therefore, the results should be interpreted as aggregate indicators rather than precise geolocation.

5.4 Interactive Results and Data Availability

To enhance transparency and facilitate the exploration of our findings, the primary results and interactive visualizations have been made available via a dedicated web dashboard[9]. This platform allows users to dynamically navigate through centrality metrics, community networks, and thematic distributions (BERTopic), providing a multi-scale view of the coordination patterns that complements the static analyses presented in this paper.

5.5 Summary of Findings and Research Contributions

The experimental results presented in this section allow us to directly address the three contributions outlined in Section 1.

[9] https://fernandanascimento26.github.io/Visualizacao-de-dados/

Contribution (i) — Operational pipeline for coordination detection on WhatsApp. The RSN-based pipeline (TF–IDF cosine similarity + 60-second sliding window, $\tau_{sim} = 0.70$, $\tau_{time} = 60$s) successfully identified coordination signals in both periods, yielding weighted networks of 170 and 237 users for PRE2022 and POS2022, respectively (Table 4). The empirical validation in Figure 2 confirms that the chosen thresholds are supported by the temporal and semantic distributions of the data, and sensitivity analyses with $\tau_{time} \in \{30\text{s}, 120\text{s}\}$ demonstrated robustness of the structural conclusions to moderate parameter variation. The pipeline provides a reproducible and platform-aligned strategy for coordination analysis in environments where interactions are not explicit.

Contribution (ii) — Characterization of coordination networks in the pre- and post-election periods. The comparison between PRE2022 and POS2022 reveals a consistent pattern of structural intensification after the electoral event. Although both networks remain sparse and fragmented, the post-election network is larger, denser in terms of coordination weight ($\bar{w} = 11.82$ vs. 5.59), and contains a higher proportion of strong edges ($w \geq 5$: 23.3% vs. 13.2%). Community detection (Louvain; $Q = 0.779$ and $Q = 0.745$) reveals a dual dynamic: a proliferation of peripheral coordination dyads alongside the consolidation of larger teams (top-3 sizes: 30, 29, 7 in POS2022 vs. 13, 13, 12 in PRE2022). Centrality analysis further confirms a functional division of labor between a densely connected internal core (high strength and eigenvector centrality) and a set of bridge actors (high betweenness), a structural signature consistent with organized amplification systems.

Contribution (iii) — Contrast between episodic and hardcore coordination. The weight-threshold analysis and

Table 13. Thematic Analysis and Frequency Variation of Terms (PRE2022 vs. POS2022). Portuguese terms are retained for reference, with English translations provided for reviewers.

Macro Category	Term (PT)	Term (EN)	Freq. PRE	Freq. POS	Abs. Diff.	% Var.	Variation
1st Round	1º turno	1st round	413	13	-400	-96.9	-97%
	primeiro turno	first round	1096	61	-1035	-94.4	-94%
Campaign	luladrão	lula-thief (slur)	270	154	-116	-43.0	-43%
	jair	Jair	1988	984	-1004	-50.5	-51%
	candidato	candidate	2166	379	-1787	-82.5	-83%
	campanha	campaign	3263	991	-2272	-69.6	-70%
	lula	Lula	21031	17559	-3472	-16.5	-17%
	bolsonaro	Bolsonaro	28174	14456	-13718	-48.7	-49%
Fraud	radios	radios	10	14	4	40.0	+40%
	roubado	stolen	99	70	-29	-29.3	-29%
	fraude	fraud	1369	1327	-42	-3.1	-3%
	inserções	ad inserts	515	31	-484	-94.0	-94%
Intervention	intervenção	intervention	103	2502	2399	2329.1	+2329%
	artigo 142	article 142	23	196	173	752.2	+752%
	art 142	art 142	14	133	119	850.0	+850%
	intervencao	intervention	4	29	25	625.0	+625%
Thief / Slur	presidiario	inmate	20	13	-7	-35.0	-35%
	ladrao	thief	190	128	-62	-32.6	-33%
	ladrão	thief	2999	2360	-639	-21.3	-21%
Stoppage / Protest	ruas	streets	927	2358	1431	154.4	+154%
	bloqueio	blockade	102	524	422	413.7	+414%
	invadir	invade	136	494	358	263.2	+263%
	greve	strike	48	398	350	729.2	+729%
	brasilia	Brasilia	26	202	176	676.9	+677%
	estrada	road	61	89	28	45.9	+46%
Polls	ibope	Ibope (pollster)	123	26	-97	-78.9	-79%
	datafolha	Datafolha	550	10	-540	-98.2	-98%
	pesquisa	poll / research	3547	257	-3290	-92.8	-93%
HQ / Barracks	quartel	barracks	15	1183	1168	7786.7	+7787%
	qg	HQ	37	1115	1078	2913.5	+2914%
	comandante	commander	89	624	535	601.1	+601%
Ballot Boxes	codigo fonte	source code	0	12	12	–	–
	biometria	biometrics	55	22	-33	-60.0	-60%
	urnas	ballot boxes	2168	2089	-79	-3.6	-4%
	urna	ballot box	554	154	-400	-72.2	-72%
Victory	vitoria	victory	40	53	13	32.5	+32%
	vencer	win	741	414	-327	-44.1	-44%
	eleito	elected	1308	567	-741	-56.7	-57%

k-shell decomposition together show that the increase in coordination intensity is not merely quantitative but structural. The proportion of users in the strict coordination core expands from 2.35% (PRE2022) to 10.97% (POS2022), and among users active in both periods, the share occupying the hard core rises from 8.3% to 19.4% (Table 6). The fraction of strong edges ($w \geq 5$) nearly doubles, and the fraction of edges at the minimum threshold ($w = 2$) decreases proportionally, indicating that recurrent coordination replaces episodic co-occurrence as the dominant pattern in the post-election ecosystem. Semantic characterization confirms that this structural consolidation is associated with politically oriented content: dominant communities in POS2022 circulate amplification calls, contestation narratives, and references to institutional intervention, while the geographic distribution of the core shifts from a diffuse multi-regional pattern to a marked concentration in Paraíba (17 core users vs. 9 in São Paulo for PRE2022), suggesting a strategic territorial reorganization of coordination efforts after the election results.

6 Conclusions

This study investigated coordination within public WhatsApp groups associated with the Brazilian far-right during

Table 14. Super-spreaders by Coordination Intensity. Primary super-spreaders by coordination intensity (W_{coord}), with capillarity (Cap.) measured as the number of groups reached by the user divided by the total number of groups in which their community was active. Truncated IDs preserve anonymity.

PRE2022					
Comm.	**User**	W_{coord}	**Rep.**	**Cap.**	**Typical message (excerpt)**
6	4ea1b3ee	260	7	3/5	visita oficial governo tocantins ministro infraestrutura gar...
6	ef496106	247	7	3/5	adolescente fere tiros três estudantes escola pública ceará ...
2	cebd2107	127	12	6/15	mundo esquerda poder venezuela fecha 15 rádios semana pingos...
2	b6242381	72	8	4/15	brasil moraes protege institutos pesquisa fraude eleitoral c...
2	19654cb7	60	8	4/9	brasil imprensa desmentida pec quer aumentar mandato senador...
POS2022					
Comm.	**User**	W_{coord}	**Rep.**	**Cap.**	**Typical message (excerpt)**
2	ef496106	764	4	2/5	mundo justiça desobedecer julgar condenar pergunte a ele ...
2	4ea1b3ee	652	4	3/6	mundo foi uma manifestação pacífica e constitucional ...
1	fac2d68f	342	3	1/1	mundo é a economia, estúpido...
2	b66e161a	261	4	2/6	mundo justa causa...
1	f9a3713e	254	3	1/1	mundo acabou a farra...

Table 15. Geographical Distribution. Approximate geographical distribution of the *expanded core* (upper layers of the k-shell), inferred from area codes (DDD) and mapped to states (UF). Area codes 0 and 99 indicate unresolved values and were omitted.

	States (core)				**Area Codes (core)**			
Rank	**PRE**	**#**	**POS**	**#**	**PRE**	**#**	**POS**	**#**
1	SP	9	PB	17	11	7	83	17
2	PB	5	SP	8	83	5	31	4
3	RJ	4	MG	6	21	4	61	4
4	RN	4	DF	4	84	4	63	4
5	PA	3	TO	4	92	2	21	4

the 2022 elections, moving beyond an aggregated structural view by operationalizing *coordination* as synchronized textual repetition within short time windows and modeling it as a weighted network between users. Using a detection engine based on *rapid spread networks* and a minimum robustness filter ($w \geq 2$), we constructed coordination networks for PRE2022 and POS2022 and analyzed how coordinated activity is organized into blocks, how it stratifies into a core–periphery structure, and how it manifests in terms of intensity, recurrence, and inter-group reach Pacheco *et al*. [2021]; Mannocci *et al*. [2024a]; Kansaon *et al*. [2024].

Our results demonstrate a significant reconfiguration in the post-election period: although both networks remain sparse and fragmented, POS2022 is larger and, above all, exhibits a higher coordination intensity. This is reflected in the increase of the total weight sum and the average weight per edge, as well as the growth in the fraction of strong connections ($w \geq 5$). These patterns indicate that, after the election, coordination occurs more repeatedly between user pairs, suggesting synchronized routines sustained by a more active set of participants (Table 4) Barabasi [2005].

The community analysis reinforces this interpretation. In the post-election period, we observed the simultaneous (i) proliferation of coordinated micro-groups and (ii) consolidation of larger *teams* that concentrate connectivity and activity, functioning as amplification hubs. In parallel, the core-periphery decomposition reveals an expansion of the *hard core* (strict core via the k-shell), with a higher fraction of users in the core in POS2022 and a concentration of the operational core within a few dominant communities. These findings support the view that the ecosystem does not merely react to the electoral event but reorganizes its amplification infrastructure toward more persistent patterns of synchronization.

In characterizing structural roles, we found a strong correlation between strength (weighted degree) and influence (eigenvector centrality), indicating that users who coordinate most often also tend to occupy influential positions within the core. Conversely, betweenness centrality remains largely independent of influence, suggesting a division of labor between internal amplification (the core) and transversal connectivity (bridges between communities). Temporally, coordination/toxicity increased in the final stretch of PRE2022, while POS2022 showed an initial peak followed by a more stable regime, maintaining a baseline of coordinated activity after the event.

From a semantic perspective, the community-level analysis highlights that structural detection is topic-agnostic and that the characterization stage is crucial for distinguishing incidental coordination from politically oriented coordination. In POS2022, dominant communities exhibit a strong amplification pattern (circulation of links/articles and calls to forward content), alongside topics associated with contestation and surveillance.

The super-dissemination analysis identifies two complementary profiles: actors with very intense coordination concentrated in a few groups, and actors with moderate intensity but high capillarity, consistent with inter-group distribution and amplification functions. Furthermore, the aggregation of misinformation scores shows a decrease in the weighted average between periods but maintains significant heterogeneity between communities, reinforcing the importance of community-level analysis. Finally, the approximate geographical distribution (by area code) suggests shifts in the relative concentration of the core across periods, though this indicator should be interpreted as an aggregate proxy.

This study has limitations inherent to its context and data. Anonymization imposes restrictions on linking narratives to identifiable actors and on the fine-grained discursive characterization of communities, although the interaction structure and metrics preserve useful signals of communicational dynamics. Additionally, geographic inference from area codes is an imperfect proxy and may reflect limitations in registration and extraction; thus, it should be interpreted with caution.

As future work, we highlight: (i) incorporating more robust semantic representations (e.g., embeddings) and modalities beyond text (audio, images, stickers) to broaden the coverage of coordination; (ii) advancing toward dynamic temporal modeling (communities over time) to measure stability, transitions, and behavioral archetypes; (iii) investigating mechanisms of causation and reactivity to external events, linking coordination time series to political milestones and inter-group diffusion; (iv) combining coordination signals with automation and infrastructure indicators (reused accounts, forwarding patterns, multi-group participation) to strengthen inferences; and (v) improving spatial inference through more reliable methods and exploring comparative analyses across different platforms.

References

Barabasi, A.-L. (2005). The origin of bursts and heavy tails in human dynamics. *Nature*, 435(7039):207–211.

Chagas, V. (2022). Whatsapp and digital astroturfing: A social network analysis of brazilian political discussion groups of bolsonaro's supporters. *International journal of communication*, 16:25.

Chagas, V. and Da-Costa, G. (2023). Whatsapp and transparency: an analysis on the effects of digital platforms' opacity in political communication research agendas in brazil. *Profesional de la información*. DOI: 10.3145/epi.2023.mar.23.

Cinelli, M., Cresci, S., Quattrociocchi, W., Tesconi, M., and Zola, P. (2022). Coordinated inauthentic behavior and information spreading on twitter. *Decision Support Systems*, 160:113819.

Cinus, F., Minici, M., Luceri, L., and Ferrara, E. (2025). Exposing cross-platform coordinated inauthentic activity in the run-up to the 2024 u.s. election. In *Proceedings of the ACM on Web Conference 2025*, WWW '25, page 541–559, New York, NY, USA. Association for Computing Machinery. DOI: 10.1145/3696410.3714698.

de Sá, I. C., Galic, L., Franco, W., Gadelha, T., Monteiro, J. M., and Machado, J. C. (2023). BATMAN: A big data platform for misinformation monitoring. In Filipe, J., Smialek, M., Brodsky, A., and Hammoudi, S., editors, *Proceedings of the 25th International Conference on Enterprise Information Systems, ICEIS 2023, Volume 1, Prague, Czech Republic, April 24-26, 2023*, pages 237–246. SCITEPRESS. DOI: 10.5220/0011995500003467.

Evangelista, R. and Bruno, F. (2019). Whatsapp and political instability in brazil: targeted messages and political radicalisation. *Internet Policy Review*, 8(4). DOI: 10.14763/2019.4.1434.

Giglietto, F., Righetti, N., Rossi, L., and Marino, G. (2020). It takes a village to manipulate the media: Coordinated link sharing behavior during 2018 and 2019 Italian elections. *Information, Communication & Society*, 23(6):867–891. DOI: 10.1080/1369118X.2020.1739732.

Hristakieva, K., Cresci, S., Da San Martino, G., Conti, M., and Nakov, P. (2022). The spread of propaganda by coordinated communities on social media. In *Proceedings of the 14th ACM web science conference 2022*, pages 191–201.

Iswarya, V., Govindasamy, V., and Akila, V. (2024). A survey on the identification of influential spreaders in complex networks. In *2024 2nd International Conference on Computer, Communication and Control (IC4)*, pages 1–4. DOI: 10.1109/IC457434.2024.10486364.

Kansaon, D., de Freitas Melo, P., Zannettou, S., and Benevenuto, F. (2025). From fake news to real protests: Whatsapp's role in brazilian political coordination. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 19, pages 1007–1020.

Kansaon, D., Melo, P. d. F., Zannettou, S., Feldmann, A., and Benevenuto, F. (2024). Strategies and attacks of digital militias in whatsapp political groups. *Proceedings of the International AAAI Conference on Web and Social Media*,

18(1):813–825. DOI: 10.1609/icwsm.v18i1.31354.

Keller, F. B., Schoch, D., Stier, S., and Yang, J. (2020). Political astroturfing on twitter: How to coordinate a disinformation campaign. *Political communication*, 37(2):256–280.

Machado, C., Kira, B., Narayanan, V., Kollanyi, B., and Howard, P. N. (2019). A study of misinformation in whatsapp groups with a focus on the brazilian presidential elections. In *Proceedings of The Web Conference (WWW '19)*. DOI: 10.1145/3308560.3316738.

Magelinski, T., Ng, L., and Carley, K. (2022). A synchronized action framework for detection of coordination on social media. *Journal of Online Trust and Safety*, 1(2). DOI: 10.54501/jots.v1i2.30.

Mannocci, L., Mazza, M., Monreale, A., Tesconi, M., and Cresci, S. (2024a). Detection and characterization of coordinated online behavior: A survey. *arXiv preprint*.

Mannocci, L., Mazza, M., Monreale, A., Tesconi, M., and Cresci, S. (2024b). Detection and characterization of coordinated online behavior: A survey. *arXiv preprint arXiv:2408.01257*.

Mudde, C. (2016). Far right parties in europe. *Annual Review of Political Science*, 19:291–308. DOI: 10.1146/annurev-polisci-042814-010551.

Mudde, C. (2019). *The Far Right Today*. Polity Press, Cambridge, UK.

Nascimento, F. F. d., Monteiro, J. M., and Machado, J. (2025). Dinâmicas de grupos de whatsapp da extrema direita no brasil: Uma análise comparativa pré e pós-eleição de 2022. In *Anais do XL Simpósio Brasileiro de Banco de Dados (SBBD 2025)*, pages 563–575, Porto Alegre. SBC. DOI: 10.5753/sbbd.2025.247285.

Newman, M. E. J. (2006). Modularity and community structure in networks. *Proceedings of the National Academy of Sciences*, 103(23):8577–8582. DOI: 10.1073/pnas.0601602103.

Nizzoli, L., Tardelli, S., Avvenuti, M., Cresci, S., and Tesconi, M. (2021). Coordinated behavior on social media in 2019 uk general election. In *Proceedings of the international AAAI conference on web and social media*, volume 15, pages 443–454.

Norris, P. and Inglehart, R. (2019). *Cultural Backlash: Trump, Brexit, and Authoritarian Populism*. Cambridge University Press.

Pacheco, D., Hui, P.-M., Torres-Lugo, C., Truong, B. T., Flammini, A., and Menczer, F. (2021). Uncovering coordinated networks on social media: Methods and case studies. In *Proceedings of the International AAAI Conference on Web and Social Media (ICWSM)*, volume 15, pages 455–466.

Pirro, A. L. P. (2022). Far right: The significance of an umbrella concept. *Nations and Nationalism*, 29(1):101–112. DOI: 10.1111/nana.12860.

Resende, G., Melo, P., Sousa, H., Messias, J., Vasconcelos, M., Almeida, J., and Benevenuto, F. (2019). (mis)information dissemination in whatsapp: Gathering, analyzing and countermeasures. DOI: 10.1145/3308558.3313688.

Rogers, R. and Righetti, N. (2025). Coordinated inauthentic behaviour on Facebook? A typology of manufactured attention. *New Media & Society*. Online first. DOI: 10.1177/29768624251369784.

Rydgren, J. (2007). The sociology of the radical right. *Annual Review of Sociology*, 33:241–262. DOI: 10.1146/annurev.soc.33.040406.131750.

Sharma, K., Zhang, Y., Ferrara, E., and Liu, Y. (2021). Identifying coordinated accounts on social media through hidden influence and group behaviours. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, KDD '21, page 1441–1451, New York, NY, USA. Association for Computing Machinery. DOI: 10.1145/3447548.3467391.

Tardelli, S., Nizzoli, L., Tesconi, M., Conti, M., Nakov, P., Martino, G. D. S., and Cresci, S. (2024). Temporal dynamics of coordinated online behavior: Stability, archetypes, and influence. *Proceedings of the National Academy of Sciences*, 121(20):e2307038121. DOI: 10.1073/pnas.2307038121.

Thiele, D., Milzner, M., Heft, A., Gong, B., and Pfetsch, B. (2025). Attributing coordinated social media manipulation: A theoretical model and typology. *Information, Communication & Society*. Online first. DOI: 10.1177/14614448251350100.

Vishnuprasad, P. S., Nogara, G., Cardoso, F., Cresci, S., Giordano, S., and Luceri, L. (2024). Tracking fringe and coordinated activity on twitter leading up to the us capitol attack. *Proceedings of the International AAAI Conference on Web and Social Media*, 18(1):1557–1570. DOI: 10.1609/icwsm.v18i1.31409.

Weber, D. and Neumann, F. (2021). Amplifying influence through coordinated behaviour in social networks. *Social Network Analysis and Mining*, 11(1):1–42. DOI: 10.1007/s13278-021-00815-2.

Zheng, H., Zhao, H., and Ahmadi, G. (2024). Towards improving community detection in complex networks using influential nodes. *Journal of Complex Networks*, 12(1):cnae001. DOI: 10.1093/comnet/cnae001.