*Journal on Interactive Systems, 2025, 16:1,*doi: 10.5753/jis.2025.5526

A Network-Driven Framework for Bidimensional Analysis of Information Dissemination on Social Media Platforms

Geovana S. Oliveira [**Federal University of Ouro Preto** | ***geovana.so@aluno.ufop.edu.br***]
João Pedro Lobo [**Federal University of Ouro Preto** | ***joao.lobo1@aluno.ufop.edu.br***]
Otávio Venâncio [**Federal University of Minas Gerais** | ***otavio.venancio@dcc.ufmg.br***]
Vinícius da F. Vieira [**Federal University of São João del-Rei** | ***vinicius@ufsj.edu.br***]
Jussara M. Almeida [**Federal University of Minas Gerais** | ***jussara@dcc.ufmg.br***]
Ana P. C. Silva [**Federal University of Minas Gerais** | ***ana.coutosilva@dcc.ufmg.br***]
Ronan S. Ferreira [**Federal University of Ouro Preto** | ***ronan.ferreira@ufop.edu.br***]
Carlos H. G. Ferreira [**Federal University of Ouro Preto** | ***chgferreira@ufop.edu.br***]

Department of Computing and Systems, Universidade Federal de Ouro Preto, João Monlevade, MG, 35931-008, Brazil.

Received: 11 February 2025 • **Accepted:** 02 September 2025 • **Published:** 17 September 2025

Abstract: Network modeling has become a foundational approach for analyzing information dissemination on social media platforms. Moreover, backbone extraction techniques, designed to isolate the most relevant structural patterns in noisy networks, have been widely employed to identify salient interactions, especially in the context of coordinated behavior and campaign detection. However, most of these approaches rely on a one-dimensional perspective, focusing primarily on interaction volume while neglecting temporal dynamics that are crucial to understanding how content spreads in real time. This limited view can obscure important distinctions between dissemination strategies that are fast but sparse or slow but voluminous. To address this gap, we introduce a bidimensional framework that integrates both interaction volume and speed, enabling a more comprehensive modeling of dissemination dynamics. Our approach applies state-of-the-art backbone extraction techniques independently to each dimension, classifying edges into four distinct dissemination profiles. This classification provides new analytical affordances for exploring both structural and textual patterns of information flow across social platforms. Applied to two case studies on Twitter/X and Telegram, the framework reveals contrasting dissemination strategies across platforms and shows how different edge classes contribute to the amplification of specific narratives. These findings advance the study of information dynamics by offering a finer-grained, multidimensional perspective on user interactions in complex social networks.

Keywords: Social Media Platforms, Information Dissemination, Network Modeling and Analysis, Network Backbone Extraction.

1 Introduction

The advent of social media platforms has profoundly transformed information dissemination [Weber and Neumann, 2021]. These platforms enable users to engage in large-scale interactions, share content instantaneously, and form virtual communities without geographical or temporal constraints [Souravlas *et al.*, 2021]. As a result, they have accelerated the spread of events, debates, and narratives, reshaping the dynamics of public discourse and political engagement [Zhuravskaya *et al.*, 2020; Stieglitz and Dang-Xuan, 2013]. This phenomenon has been widely documented across multiple geopolitical contexts. Examples include the 2016 presidential election and the Capitol riots in January 2021 in the United States [Vishnuprasad *et al.*, 2024; Bizel and Singh, 2023; Golovchenko *et al.*, 2020; Badawy *et al.*, 2019]. Similarly, in Brazil, the 2022 presidential elections and the January 8, 2023 riots illustrate the role of social media in shaping political events [Venâncio *et al.*, 2024a; Rossini *et al.*, 2023].

One of the most widely adopted approaches for analyzing and understanding information dissemination in social networks is graph-based modeling [Lima *et al.*, 2024; Antonakaki *et al.*, 2021; Zayats and Ostendorf, 2018]. This approach represents user interactions as a structured network composed of nodes (individuals or accounts) and edges (relationships between them) [Barabási, 2013]. Due to its effectiveness in identifying content dissemination patterns and supporting analysis at different levels of granularity, graph modeling has been extensively explored in the literature [Rossetti *et al.*, 2017; Abilov *et al.*, 2021].

Despite the advantages of using social networks to study information diffusion and interaction dynamics, these models often exhibit high connectivity, with many weak and sporadic links that may not be critical to understanding the phenomenon under investigation. Such transient interactions can introduce noise into the analysis, obscuring meaningful patterns [Gomes Ferreira *et al.*, 2022; Neal *et al.*, 2021]. To address this issue, backbone extraction methods have become essential, as they filter out less relevant connections and highlight the most significant and recurring relationships within a network [Pacheco *et al.*, 2021, 2020; Nobre *et al.*, 2020; Ferreira *et al.*, 2020; Nobre *et al.*, 2022]. By retaining only the most relevant edges for a given study, these methods produce a simplified yet more representative structure—

commonly called the *network backbone* [Coscia and Neffke, 2017]. Various techniques are available for backbone extraction, each tailored to different objectives and applications [Gomes Ferreira *et al.*, 2022; Neal *et al.*, 2021]. These methods are frequently employed in studies on collective information dissemination in social media, providing more precise insights into network structures and user interactions [Neal *et al.*, 2021; Coscia and Neffke, 2017; Linhares *et al.*, 2022; Araujo *et al.*, 2023].

Although graph-based modeling and backbone extraction techniques have been successfully applied to study information diffusion on social media, their implementation remains predominantly one-dimensional [Nobre *et al.*, 2020; Linhares *et al.*, 2022; Venâncio *et al.*, 2024a; Pacheco *et al.*, 2020; Weber and Neumann, 2021; da Rosa *et al.*, 2022; Ferreira *et al.*, 2021]. In most studies, network structures are defined solely based on the volume of user interactions, overlooking other critical aspects of information dissemination. For instance, it is common to model networks by considering the number of co-retweets shared between users on Twitter/X [Linhares *et al.*, 2022; Pacheco *et al.*, 2021] or the volume of similar messages exchanged among participants in WhatsApp and Telegram groups [Nobre *et al.*, 2020, 2022; Venâncio *et al.*, 2024a].

While extracting the backbone based on interaction volume allows for the characterization of network topology and the identification of community structures – often revealing potential patterns of coordination – this approach neglects essential factors that influence the spread of information. One of the most critical yet frequently overlooked aspects is the *speed* at which information propagates. The velocity of content dissemination plays a fundamental role in shaping the dynamics and impact of information diffusion, influencing the reach, engagement, and potential virality of specific narratives [Varshney and Vishwakarma, 2021; Lanius *et al.*, 2021; Elmas *et al.*, 2021; Bellutta and Carley, 2023]. By relying solely on interaction volume, existing models may fail to capture key mechanisms underlying the amplification of content and the emergence of rapidly spreading information cascades, limiting the depth of analysis regarding the influence of different dissemination strategies.

In this paper, we simultaneously address the volume and speed of interactions on social media platforms and analyze how these two dimensions influence content dissemination through network-based models. To this end, we build upon our previous framework [Oliveira *et al.*, 2024], extending its flexibility and applicability to better characterize multiple facets of information dissemination on online social networks. This extension enhances the analytical capacity of existing approaches by integrating structural network modeling with textual pattern mining, allowing for a more comprehensive examination of information dissemination dynamics. Our investigation focuses on how the topological patterns of user interactions affect content propagation, explicitly considering the interplay between interaction volume and temporal dynamics. By incorporating textual analysis, we bridge a critical gap in the literature, associating interaction pattern classification with the thematic structure of the information being propagated. The central hypothesis underlying this study is that different bidimensional dissemination patterns (i.e., varying interaction volume and speed) play distinct roles in shaping the reach and prominence of specific content. Building on this hypothesis, we seek to answer the following research questions:

RQ1: How can we identify and characterize groups of users with distinct information dissemination patterns in terms of volume and interaction speed on social media platforms? To address this, we model co-interaction networks, separating them into two dimensions: volume and speed of interactions. We then apply backbone extraction techniques to identify the most relevant edges in these networks and classify them into different dissemination profiles. These profiles include four categories: (1) users who disseminate a low volume of information at a slow pace; (2) users who disseminate a high volume of information but at a slow pace; (3) users who disseminate a low volume of information but at a fast pace; and (4) users who disseminate a high volume of information at a fast pace. Subsequently, we analyze the topology of these networks and the organization of users into communities to understand dissemination patterns.

RQ2: To what extent do bidimensional dissemination patterns (volume and speed), represented by edge classifications, contribute to the dissemination of specific topics on social media platforms? To answer this question, we integrate topic analysis into our network modeling approach using BERTopic. This technique enables us to identify and track the spread of topics within the network, linking them to the dissemination profiles identified in RQ1. By correlating topic emergence with distinct interaction patterns, we evaluate whether certain types of content are preferentially amplified by users who engage in rapid dissemination, high-volume dissemination, or a combination of both. This approach allows us to assess the influence of dissemination profiles on shaping discursive landscapes and narrative propagation within social media ecosystems.

We apply our framework to two case studies on Twitter/X and Telegram, focusing on periods with strong evidence of coordinated campaigns aimed at spreading misinformation or influencing public opinion [Araujo *et al.*, 2023; G. Da Fonseca *et al.*, 2024; Venâncio *et al.*, 2024a]. Our findings demonstrate that the proposed framework enables a more nuanced understanding of information dissemination dynamics by identifying distinct profiles based on interaction volume and speed (RQ1). On Twitter/X, dissemination is characterized by more cohesive communities and high-volume interactions, reflecting widespread but coordinated content sharing. In contrast, Telegram exhibits more fragmented yet ideologically aligned structures, where slower, lower-volume dissemination prevails, suggesting a platform more oriented toward strategic mobilization than rapid amplification. Furthermore, by integrating topic modeling (RQ2), we show that different dissemination patterns play distinct roles in amplifying specific narratives. While some user groups predominantly reinforce dominant themes, others act as accelerators of information spread, revealing the interplay between dissemination structures and the thematic composition of the content.

Our main contributions are:

- Establishment of a bidimensional framework for analyzing information dissemination. The proposed framework integrates state-of-the-art backbone extraction methods. It offers a flexible approach applicable across different social media platforms, as it only requires data on interaction volume and speed. This flexibility ensures adaptability to various online ecosystems and methodological setups.
- Identification of multiple dissemination patterns through structural and textual pattern mining. Unlike traditional models focusing solely on interaction structures, our framework enables the integration of diverse textual analysis techniques, including topic modeling, sentiment analysis, and toxicity detection—allowing a broader characterization of amplified narratives and user engagement strategies.
- Empirical validation through two case studies on Twitter/X and Telegram. The application of our framework reveals distinct propagation patterns across different platforms, demonstrating how structural differences in user interactions influence content dissemination and mobilization strategies. Our findings highlight the versatility of the approach in uncovering both platform-specific and generalizable dissemination behaviors.

The remainder of the work is organized as follows: Section 2 presents the main fundamentals on which our framework is based and related work. Section 3 details the framework, while Section 4 describes the case studies. Then, Section 5 presents the results, and Section 6 brings the conclusions and future directions.

2 Background

This section presents the fundamentals and previous studies related to our work. First, we discuss fundamental concepts in network modeling, particularly community detection and backbone extraction techniques. Then, we review related studies on information dissemination in social media, emphasizing coordinated behaviors and the challenges of analyzing large-scale interaction networks.

2.1 Fundamentals

Network modeling has been widely used to study both online and offline phenomena, including the dissemination of information on digital platforms [Newman, 2006; Fortunato, 2010; Barabási, 2013]. One of the main approaches for analyzing such networks is community detection, which seeks to identify subsets of nodes that exhibit denser interactions among themselves than with the rest of the network [Khan and Niazi, 2017; Rossetti *et al*., 2017]. This structuring enables a better understanding of user organization and information dissemination patterns across various contexts, such as social networks, biological systems, and communication infrastructures. The central idea behind community detection is that nodes within the same community share common characteristics and play similar roles in information propagation. Well-defined communities can indicate interest groups, informational bubbles, or even coordinated disinformation networks, depending on the phenomenon under study. However, identifying these communities efficiently is a computationally complex challenge, as the problem is NP-hard, meaning no exact algorithms can solve it in polynomial time for large-scale networks [Fortunato and Hric, 2016].

Given this complexity, the most commonly used methods for community detection are heuristics based on modularity optimization, a metric that measures the quality of a network partition into communities [Newman and Girvan, 2004]. Modularity compares the density of edges within identified communities with the expected density in a null model, maximizing the internal cohesion of groups. Among the most popular algorithms utilizing this approach are Louvain [Blondel *et al*., 2008] and Leiden [Traag *et al*., 2019], both widely used in the literature due to their efficiency and scalability for large networks.

As the amount of data and, consequently, interactions on these platforms increases, a significant challenge in network analysis arises: the presence of noise in connections. Social network topologies can be excessively dense, with weak or sporadic edges that may obscure relevant patterns. Including all edges indiscriminately in network modeling can lead to misinterpretations and hinder the detection of meaningful communities [Coscia and Neffke, 2017; Newman, 2018]. Backbone extraction techniques have emerged to mitigate this issue, aiming to filter out less relevant connections and produce a reduced yet representative version of the network. This process enables the capture of the most essential structures and facilitates understanding underlying patterns. The concept of a *network backbone* refers to a simplified version of the original network in which only the most relevant connections are retained. The importance of an edge can be defined based on different criteria, considering either local node characteristics or global network properties [Markopoulou *et al*., 2008; Oliveira *et al*., 2024; Neal, 2022]. Backbone extraction methods can be categorized into different classes, depending on the approach used to determine edge relevance.

Several studies have focused on evaluating the properties of these methods [Gomes Ferreira *et al*., 2022; Neal and Neal, 2024; Neal, 2022]. One key distinction among these methods is how edge salience is determined—locally or globally. Local methods assess the importance of connections based on the immediate context of nodes and their interactions, allowing for the capture of specific nuances in certain network regions. On the other hand, global methods apply uniform criteria to all edges, considering overall network structures. While local methods are helpful in highlighting specific interaction patterns, global methods provide a panoramic view of the most significant connections across the network.

Moreover, backbone extraction methods can be classified into structural and probabilistic approaches. Structural methods rely exclusively on topological attributes, such as edge weights, neighborhood overlap, and path definitions, to determine connection relevance. Examples of this category include the *High Salient Skeleton* [Grady *et al*., 2012], which removes edges based on structural criteria, and *Thresholding*, which establishes a minimum threshold for connection inclusion. Probabilistic methods, in contrast, use statistical reference models to evaluate edge relevance by comparing them to expected distributions in a null model. For instance, the *Disparity Filter* analyzes the distribution of edge weights to identify those that significantly deviate from a random model, enabling the detection of more relevant connections [Serrano *et al.*, 2009]. The *Polya Urn Filter* extends this method by incorporating a more refined model to capture local heterogeneities in interactions [Marcaccioli and Livan, 2019].

Specifically, the *Polya Urn Filter* is based on the principle of stochastic reinforcement, where the probability of a connection being strengthened over time is proportional to its frequency of occurrence. This model captures the preferential attachment dynamics observed in social networks, where recurring interactions tend to consolidate stronger ties among users [Gomes Ferreira *et al.*, 2022]. A reference model is built for each edge, capable of capturing the reinforcement of existing interactions by examining the degree and strength (sum of the weights of all edges incident to the node) of each node connected to that edge. This reinforcement mechanism can be regulated and fine-tuned. Connections deviating significantly from expected behavior (given by the reference model) are considered salient and retained in the *backbone* [Marcaccioli and Livan, 2019].

The main advantage of the *Polya Urn Filter* over other approaches is its ability to respect the intrinsic heterogeneity of social networks [Gomes Ferreira *et al.*, 2022]. Instead of applying a fixed criterion to all edges, the method adapts its evaluation based on the local structure of each node, ensuring that relevant connections across different scales are properly identified. This approach is particularly useful for studying highly dynamic and imbalanced networks, where dissemination patterns can vary significantly across different subgroups. In this work, we used the *Polya Urn Filter* due to its ability to capture heterogeneous information dissemination patterns and its suitability for modeling complex social networks. By combining a statistical approach with local characteristics, this method allows for a more representative backbone extraction, contributing to a more accurate analysis of information propagation mechanisms on digital platforms. However, other methods could also be used, as discussed in the methodology section.

2.2 Related Work

Network modeling has been extensively used to detect and analyze information dissemination in social media platforms such as Twitter/X, Telegram, WhatsApp, and Instagram [Keller *et al.*, 2020; Dargahi Nobari *et al.*, 2017; Urman and Katz, 2020; Cinelli *et al.*, 2022; Weber and Neumann, 2020; Sharma *et al.*, 2021; Graham *et al.*, 2024; Giglietto *et al.*, 2020; Vargas *et al.*, 2020; Barbosa *et al.*, 2019]. Analyzing these networks enables a deeper understanding of user interaction patterns and how certain content spreads, aiding in the identification of possible coordinated campaigns, viral phenomena, and the formation of communities with shared interests. Within this scope, various approaches have been developed to model the structure of these networks and extract meaningful insights regarding information propagation mechanisms.

A significant body of research focuses on the dissemination of political information based on data extracted from social media, including Facebook [Giglietto *et al.*, 2020], WhatsApp [Nobre *et al.*, 2022], Instagram [Ferreira *et al.*, 2021], and Twitter/X [da Rosa *et al.*, 2022]. These studies seek to understand how user groups interact to amplify specific narratives, often employing coordinated strategies to influence public debate.

Among these efforts, several investigations have identified and characterized coordinated actions on Twitter/X [Pacheco *et al.*, 2020; Keller *et al.*, 2020; Grimminger and Klinger, 2021]. For instance, Pacheco *et al.* [2020] examined disinformation campaigns targeting the Syrian Civil Defense on Twitter by modeling connections between users who recurrently shared similar tweets. This behavior suggested a coordinated effort to amplify particular narratives. Similarly, Keller *et al.* [2020] analyzed retweet networks to investigate the coordination of political advertisements during the South Korean presidential election, while Danaditya *et al.* [2022] characterized synchronized interactions in political discourse on Twitter/X related to the Indonesian elections.

Most of these studies rely on networks representing direct user interactions, such as retweets, mentions, and hashtag sharing, without employing filtering techniques to distinguish structurally relevant connections from potentially noisy ones. As a result, such models may capture sporadic and superficial interactions, leading to biased interpretations of group organization and information dissemination within these platforms. Given the challenge of modeling dense and noisy networks, an increasing number of studies aim to identify and remove connections that may represent random interactions, refining the network structure to retain only the most salient edges [Nobre *et al.*, 2022, 2020; Pacheco *et al.*, 2020, 2021; Nizzoli *et al.*, 2021; da Rosa *et al.*, 2022; Ferreira *et al.*, 2019, 2020, 2021; Venâncio *et al.*, 2024a]. This process results in the network's *backbone*. This substructure preserves only the most significant connections for the studied phenomenon, facilitating the identification of relevant structural patterns in content dissemination.

For example, Pacheco *et al.* [2020] investigated disinformation campaigns against the Syrian Civil Defense on Twitter/X using a backbone extraction technique based on *thresholding*, eliminating low-weight connections to highlight persistent interaction patterns. In [Pacheco *et al.*, 2021], the same authors applied a similar approach to analyze coordinated actions in different scenarios, emphasizing the importance of filtering to remove noise and focus on the most significant interactions. Beyond *thresholding*, other techniques have been applied to refine networks modeled from social interactions. Linhares *et al.* [2022] and da Rosa *et al.* [2022] analyzed user communities involved in spreading electoral fraud claims in the U.S. 2020 elections, employing the *Disparity Filter* [Serrano *et al.*, 2009] and *Tripartite Backbone Extraction (TriBE)* methods [Ferreira *et al.*, 2021; Gomes Ferreira *et al.*, 2022]. Similarly, Araujo *et al.* [2023] investigated the coordinated promotion of early political campaigns on Twitter/X during the 2022 Brazilian preelection period, applying the *Disparity Filter + Neighborhood Overlapping* to detect salient interactions.

Other studies have explored using backbone extraction to analyze media-sharing networks. Nobre *et al.* [2020] applied the *Disparity Filter* to study co-sharing networks. In contrast, Venâncio *et al*. [2024a] examined coordination patterns on Telegram during the 2022 Brazilian elections and the January 8 political events, using a combination of the *Disparity Filter* and neighborhood overlap techniques. These previous studies have thus shown that different user groups may adopt varied strategies to disseminate information [Lanius *et al*., 2021; Varshney and Vishwakarma, 2021]. Some groups operate in a highly coordinated manner, rapidly pushing large volumes of content, while others act more diffusely, injecting sporadic messages to hinder the detection of organized patterns. Studies also indicate using temporary and paid accounts to accelerate these processes [Albarrak *et al*., 2020].

Complementing network-based approaches, an important branch of research in many such efforts has focused on text analysis to capture the content and linguistic features of discourse in social media. These studies often investigate sentiment polarity, toxicity, topic evolution, and psycholinguistic traits in user-generated messages. Tools like SentiStrength have been widely applied to infer sentiment in short texts [Ferreira *et al*., 2021; Araujo *et al*., 2023], while others employ lexical-based resources such as LIWC and LeIA for deeper psycholinguistic insight [Barbosa *et al*., 2019; da Rosa *et al*., 2022; Danaditya *et al*., 2022]. For instance, Ferreira *et al*. [2021] analyzed the emotional tone of political comments on Instagram, combining sentiment and psycholinguistic features, and Araujo *et al*. [2023] assessed the toxicity of political retweets using Google's Perspective API, highlighting the role of negative discourse in early campaigns. Similarly, Barbosa *et al*. [2019] applied LeIA to evaluate sentiment trends in tweets about social reforms, identifying emotionally charged reactions directed at political institutions.

Textual features have also been used to characterize coordinated behavior or manipulative communication strategies. Danaditya *et al*. [2022] integrated psycholinguistic markers and topic themes to identify coordinated agents in political discourse. Topic modeling is another central component of textual analysis in this domain. Traditional methods such as TF-IDF and LDA have been employed to summarize dominant themes and identify community-specific vocabularies [Barbosa *et al*., 2019; Ferreira *et al*., 2021; Nobre *et al*., 2022; Araujo *et al*., 2023]. For example, Nobre *et al*. [2022] used TF-IDF and bag-of-words representations to compare WhatsApp messages with pre-verified disinformation claims, while Ferreira *et al*. [2021] applied TF-IDF to characterize content shared by ideological clusters. Keller *et al*. [2020] analyzed word frequencies and topic shifts in astroturfing campaigns during political events, identifying lexical distinctions between coordinated actors and regular users. Albarrak *et al*. [2020] tracked financial keywords and tone in economic tweets to detect speculative behaviors.

More recent works have adopted embedding-based topic modeling methods such as BERTopic, which combine Sentence-BERT representations with dimensionality reduction and clustering (e.g., UMAP + HDBSCAN) to extract semantically rich and temporally dynamic themes [Venâncio *et al*., 2024a; da Rosa *et al*., 2022]. Venâncio *et al*. [2024a] used this approach to study Telegram discussions around military intervention and political unrest in Brazil, while da Rosa *et al*. [2022] analyzed how topic salience evolved in electoral contexts and how specific narratives were amplified by coordinated users. Additionally, simpler classification-based strategies have been applied to label and quantify content. For instance, Dargahi Nobari *et al*. [2017] used keyword filtering and content categorization to detect advertisements and spam in Telegram channels. Likewise, Albarrak *et al*. [2020] analyzed financial discourse using domain-specific lexicons and tonal classification.

Despite these advancements, most studies utilizing backbone extraction focus solely on interaction volume while overlooking a critical factor, notable, the speed at which information spreads. Moreover, while several works incorporate textual analysis, whether for sentiment, toxicity, or topic modeling, they often do so independently of the underlying network dynamics. This separation limits our understanding of how structural interaction patterns shape, constrain, or amplify the diffusion of specific narratives. Thus, Our work bridges this gap by proposing a bidimensional framework that jointly considers volume and speed of interaction to characterize dissemination profiles. By correlating these structural dimensions with textual signals, our approach enables the identification of distinct linguistic and thematic patterns associated with different dissemination behaviors such as coordinated high-speed amplification or slow, high-volume diffusion. This integration opens space for richer interpretations of content propagation, revealing how network structure and textual content co-evolve in social media environments.

To summarize, this work builds on an increasing shift in literature towards multidimensional frameworks of analysis that are able to capture who interacts with whom and how and what is communicated. By unifying structural and textual perspectives, our framework offers a more nuanced and interpretable view of information dissemination across platforms.

3 A Bidimensional Framework for Information Dissemination

Figure 1 illustrates the structure of our framework. Initially, we construct networks that represent user interactions, assigning weights that capture the volume and speed of these interactions. These networks can be built using various metrics, such as the number of *retweets*, the frequency of mentions or shared links, and the average time between consecutive interactions. This flexibility allows for a more nuanced analysis of information dissemination across different social media platforms. We then apply *backbone* extraction as described in Section 2 in each network.

Once the *backbone* is extracted, we classify edges into four distinct dissemination profiles based on their interaction volume and speed. This classification enables us to differentiate between patterns that range from sporadic and slow interactions to high-volume, fast dissemination. Leveraging this categorization, our framework enables the extraction of structural and textual patterns. Structurally, it allows the identification of communities of users with similar dissemination behaviors, centrality measures, and topological features. Beyond structural analysis, it captures textual patterns within

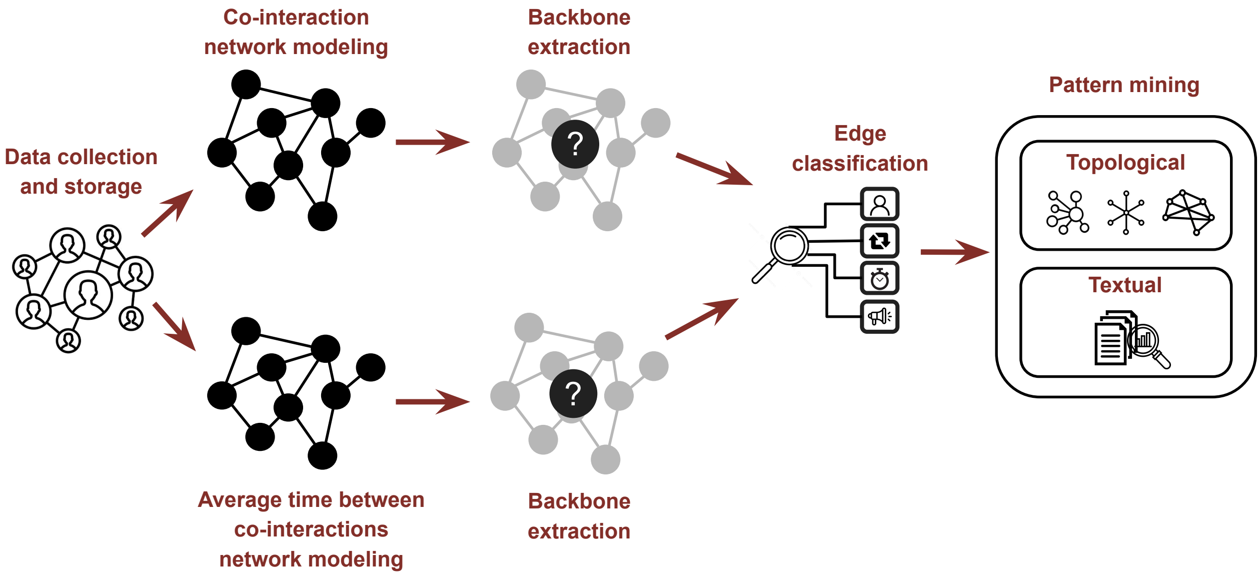

Figure 1. Overview of the proposed *framework*.

different dissemination classes, such as sentiment polarity, toxicity levels, stance on specific topics, and thematic evolution through topic modeling. Crucially, our approach is agnostic to platform, methods used for backbone extraction as well as for structural and textual analysis, allowing it to be applied across diverse social media environments. This flexibility ensures its relevance across various online ecosystems, enhancing its applicability for analyzing information diffusion in different contexts. All code implementations used in this study, including data preprocessing, network modeling, and edge classification, are publicly available in a dedicated GitHub repository[1] to support reproducibility and transparency. The following sections provide a detailed breakdown of each component of our framework.

3.1 Network Modeling

The methodology proposed in this work relies on the construction of two complementary network models: one that captures the volume of co-interactions between users and another that represents the time interval between these co-interactions within a given period of interest. We define a *co-interaction* as an event in which two or more users engage with the same or nearly identical content on a social media platform. This includes actions such as liking, sharing, commenting on, or viewing the same post, as well as publishing identical or highly similar messages, a behavior often associated with coordinated activity or shared narrative engagement Nobre *et al.* [2020, 2022]; Pacheco *et al.* [2020, 2021]. Notably, co-interactions do not imply direct interpersonal communication between users, but rather, they reflect a convergence of attention or intent around specific pieces of content. This definition allows the network model to capture structural traces of collective behavior that may arise organically or through coordinated efforts, even in the absence of explicit user-to-user links. Network-based methodologies are particularly well suited to analyze these patterns, offering insights into how information propagates, how narratives are reinforced, and how communities engage with content in distributed environments.

The co-interaction network is represented by a graph $G_{\text{volume}} = (V, E_{\text{volume}})$, where each element $v_i \in V$ represents a user, and each element $e_{ij} \in E_{\text{volume}}$ is an edge connecting a pair of users. The edge weight $w_{\text{volume}}(e_{ij})$ represents the number of co-interactions between users v_i and v_j.

In the second network, $G_{\text{time}} = (V, E_{\text{time}})$, the edges $e_{ij} \in E_{\text{time}}$ store information about the elapsed time between interactions between users represented by v_i and v_j. While other functions can be used to aggregate the times of all co-interactions between v_i and v_j, in this work, a weight $w'_{\text{time}}(e_{ij})$ is assigned to the relationship between each pair of vertices v_i and v_j and refers to the average co-interaction time between the users represented by v_i and v_j. Equation 1 defines the average time between co-interactions for each pair of users.

$$w'_{\text{time}}(e_{ij}) = \left\lceil \frac{1}{w_{\text{volume}}(e_{ij})} \sum_{p=1}^{w_{\text{volume}}(e_{ij})} t_{ij}^{p} \right\rceil, \tag{1}$$

where t_{ij}^{p} denotes the elapsed time in the p-th interaction between v_i and v_j.

The algorithms used in this work (both for backbone extraction and community detection) assume that the weight of an edge e_{ij} should be directly proportional to the strength of the relationship between v_i and v_j. However, this is counterintuitive to the definition of the weight function w'_{time} previously presented. For instance, a short time between the

[1] `https://github.com/gseovana/bidimensional-network-backbone-framework`

actions of v_i and v_j in sending the same message on a social network should indicate a strong relationship between v_i and v_j, but the definition of w'_{time} does not reflect this. To align the weight function with an intuitive interpretation, weights are normalized according to Equation 2, ensuring that higher weight values represent stronger relationships. The normalized weight $w_{\text{time}}(e_{ij})$ is given by:

$$w_{\text{time}}(e_{ij}) = \max(W_{\text{time}}) - w'_{\text{time}}(e_{ij}) + 1, \qquad (2)$$

where $\max(W_{\text{time}})$ is the highest weight found in the network G_{time}. The addition of $+1$ ensures that all resulting edge weights are strictly positive integers. This adjustment is critical for compatibility with several backbone extraction methods, such as the Polya Urn Filter, Disparity Filter, and TriBE, which assume non-zero, typically positive integer weights to model edge strength and statistical significance Ferreira *et al.* [2021]. By ensuring that the minimum possible weight is 1, we avoid eliminating potentially relevant edges and maintain consistency in the treatment of sparse temporal interactions. This transformation can be seen as a practical approximation that preserves edge relevance while ensuring compatibility with methods that require strictly positive integer weights.

In this work, each edge e_{ij} in the network modeled by G can be associated to the co-interaction of v_i and v_j to a specific content and it should be noted that the proposed methodology allows different definitions for network modeling, depending on the context and the phenomenon under study. In addition to considering the accumulated number of interactions between pairs of users, as represented in the network G_{volume}, variations can be introduced to account for the number of interactions within a critical time window. This approach is particularly useful in scenarios where dissemination behavior is highly dynamic, such as during elections, social crises, or coordinated campaigns. Similarly, the network G_{time} can be adapted to incorporate alternative metrics beyond the average interaction time. For instance, the time between consecutive interactions involving the same pair of users such as replies, message forwards, or repeated reactions can capture direct engagement patterns in both public and private settings. The time until the first response is relevant in contexts such as automated customer service, technical support, or disinformation campaigns that rely on rapid reactions to events. Additionally, the standard deviation of interaction times helps identify variations in communication pace, highlighting interactions that alternate between periods of high and low activity. The framework's flexibility allows these metrics to be adapted to different platforms and contexts, making it applicable to various information dissemination studies independent of the analyzed social media platform.

3.2 Backbone Extraction

Based on the representation of user interactions considering the bidimensional network model—taking into account both volume and speed of interaction. The framework includes a step for extracting the backbones of both networks. Our main argument in this step is the hypothesis that some edges represent strong and consistent co-interactions over the observed period, while others are sporadic, occurring in exceptionally short or long intervals. Thus, we analyze the phenomenon from the perspective of both the co-interaction volume and the time interval in which they appear. To illustrate this, consider the following scenarios on a social media platform: (1) a message is massively disseminated, shared by many users who are not very active on the platform, forming co-interaction edges with low weights; and (2) another set of messages is consistently and repeatedly disseminated by another group of users, forming edges with higher weights. Similarly, we can analyze the temporal aspect of message dissemination by determining whether interactions occur rapidly or over extended periods. Both perspectives, co-interaction volume and the temporal dynamics of these interactions, have been independently recognized in the literature as critical to understanding coordinated information dissemination Varshney and Vishwakarma [2021]; Lanius *et al.* [2021]; Ferreira *et al.* [2021]; Venâncio *et al.* [2024a]. However, they have traditionally been analyzed in isolation. The combined incorporation of these two dimensions into the analysis of information dissemination represents a novel contribution of the proposed framework.

Towards identifying the relevant edges, we propose using the *Polya Urn Filter* method [Marcaccioli and Livan, 2019] to extract the backbones of both co-interaction networks. We choose this approach, over all existing backbone extraction methods, due to its flexibility and ability to account for the local importance of edges in backbone identification, as discussed in Section 2. *Polya Urn* is applied to graph G_{volume}, generating the backbone B_{volume}, which captures nodes and edges with higher-than-expected co-interaction patterns. Similarly, the backbone extraction method must be applied to graph G_{time}. It is important to note that the weights w_{time} are normalized according to Equation 2 to ensure that the backbone extraction functions as intended, capturing edges with higher-than-expected weights [Marcaccioli and Livan, 2019; Gomes Ferreira *et al.*, 2022]. Thus, applying the *Polya Urn* method to graph G_{time} generates the backbone B_{time}, which captures edges with dissemination times that significantly deviate from expectations. By extracting the backbones of both networks, G_{volume} and G_{time}, we can classify the retained edges as representative of strong and consistent interactions. In contrast, removed edges represent sporadic and less significant interactions. This allows us to extract from the network the most important nodes and edges in terms of both information volume and dissemination time, which are classified as explained in the next section.

It is important to note that, in previous studies of information dissemination on social media, backbone extraction typically relies on a single network dimension, most often, the volume of interactions, as the sole indicator of relevance [Ferreira *et al.*, 2021; Nobre *et al.*, 2022; Venâncio *et al.*, 2024a]. In such cases, edges that exhibit high co-interaction volume (i.e., those retained in B_{volume}) are considered the structural core of the network, serving as the basis for identifying coordinated behavior or influential actors. Within our framework, these correspond to edges from both Class 2 and Class 4 that are retained in the volume-based backbone, regardless of their temporal dynamics. By incorporating a second, orthogonal dimension, i.e., interaction speed, we can uncover additional patterns not captured by volume alone. This allows the identification of edges that, while not frequent, occur in tightly clustered time windows (Classes 3 and 4), highlighting bursts of rapid dissemination that may signal coordinated amplification efforts. Conversely, our framework also identifies edges that are frequent but temporally dispersed (Class 2), capturing consistent, long-term reinforcement of narratives. Thus, the proposed bidimensional approach generalizes traditional single-backbone methods and extends their analytical scope by allowing the detection of both sustained and rapid dissemination strategies. This expanded perspective offers a more comprehensive understanding of the mechanisms underlying information spread, particularly in politically charged or coordinated contexts. The classification and comparative analysis of edge classes detailed in the next section, further highlight the interpretive value of this approach.

3.3 Edge Classification

Both network models considered in the proposed methodology (represented by G_{volume} and G_{time}) correspond to the same set of edges, associating each pair of vertices that represent users who co-interacted on social networks. Thus, it is possible to define four edge classes based on their presence or absence in B_{volume} and B_{time}, the backbones extracted from G_{volume} and G_{time}, respectively. Let's define E^B_{volume} and E^B_{time} as the sets of edges that are retained in the backbones B_{volume} and B_{time}, respectively. We can then formally define each edge class as:

- `Class 1 (low volume and low speed)`:

$$C1 = \left(E_{\text{volume}} - E^B_{\text{volume}}\right) \cap \left(E_{\text{time}} - E^B_{\text{time}}\right) \tag{3}$$

  These edges are not retained in any of the backbones.
- `Class 2 (high volume and low speed)`:

$$C2 = E^B_{\text{volume}} \cap \left(E_{\text{time}} - E^B_{\text{time}}\right) \tag{4}$$

  These edges are retained in B_{volume} but not in B_{time}.
- `Class 3 (low volume and high speed)`:

$$C3 = \left(E_{\text{volume}} - E^B_{\text{volume}}\right) \cap E^B_{\text{time}} \tag{5}$$

  These edges are retained in B_{time} but not in B_{volume}.
- `Class 4 (high volume and high speed)`:

$$C4 = E^B_{\text{volume}} \cap E^B_{\text{time}} \tag{6}$$

  These edges are retained in both backbones.

As an illustration, in the context of Twitter/X, the classified edges can be interpreted as follows: `Class 1` indicates that two users retweeted a small volume of tweets in common and with a long time interval between such tweets, on average, suggesting low volume and low dissemination speed. `Class 2` indicates that two users retweeted many tweets in common, but with long time interval between them, suggesting high volume but low dissemination speed. `Class 3` indicates that two users retweeted only a few tweets in common within a very short time interval, suggesting high dissemination speed but low volume. `Class 4` indicates that two users retweeted a high volume of tweets in common within a short time interval, suggesting strong coordination between them. In summary, the proposed approach differs from existing methodologies in the literature by enabling the differentiation of profiles and structures beyond a one-dimensional perspective.

It is worth noting that by explicitly separating volume and speed dimensions, our approach generalizes and extends previous methods of backbone extraction, offering a more fine-grained and interpretable structure for analyzing information dissemination. While the integration of alternative backbone or community detection techniques is feasible within our modular design, such expansions are better suited for follow-up studies, given the absence of standardized ground truth and the exploratory nature of unsupervised network analysis tasks.

Importantly, we highlight that `Classes 2` and `4`, which correspond to edges retained in the volume-based backbone—are aligned with the types of connections typically captured by traditional backbone extraction methods that rely on co-interaction frequency alone. In this sense, the union of `Classes 2` and `4` serves as a natural baseline for comparison with prior work, which adopts a one-dimensional view of network importance. What our framework offers, however, is a broader analytical lens. By explicitly including `Classes 1` and `3`, we are able to detect edges and patterns of dissemination that would be entirely excluded from single-criterion approaches, such as low-volume but fast amplification, or slow and persistent propagation. This classification structure allows us to compare dissemination dynamics both within and beyond the conventional backbones, reinforcing the value of adopting a bidimensional perspective.

3.4 Extraction of Structural and Textual Dissemination Patterns

The final stage of our framework consists of analyzing the generated networks, enabling the identification of structural and textual patterns that characterize information dissemination. This approach is entirely agnostic regarding the analyzed platform, allowing its application to different contexts and types of interactions. A wide range of techniques can be employed to study the structural properties of networks, such as centrality measures, identification of influential nodes, network resilience analysis, and extraction of substructures specialized in content propagation. From a textual perspective, aspects such as sentiment analysis, detection of polarized narratives, identification of linguistic patterns associated with disinformation, and evaluation of the influence of specific topics on dissemination dynamics can be explored.

Beyond these perspectives, other dimensions can be incorporated into the analysis. For example, content toxicity can be examined to determine whether specific information dissemination patterns are associated with aggressive or harmful discourse. Similarly, detecting coordinated behavior in spreading misleading or manipulated information can help identify disinformation campaigns. Integrating these different layers—structural, temporal, and textual—enhances the understanding of how information spreads and evolves in online environments.

This work focuses on community detection, aiming to understand how users organize themselves and form cohesive dissemination structures with different interaction patterns. Additionally, we perform topic modeling, investigating the relationship between network structure and the content being propagated. This allows us to analyze how different dissemination profiles influence the amplification of specific themes.

3.4.1 Community Detection

Considering the edge classification described in the previous section, four subnetworks are identified, each corresponding to one of the edge classes based on interaction volume and speed. By examining these subnetworks, we define dissemination profiles by analyzing the observed patterns within different classes, investigating the topological organization that emerges from each subnetwork, and exploring their similarities and differences. In particular, we focus on how nodes in these subnetworks are organized into community structures—i.e., subsets of nodes with denser internal connections than the rest of the network. Community organization analysis has been successfully applied to social media platforms to identify ideological groups, characterize early political campaigns, and detect misinformation or fake news [Ferreira *et al*., 2020; Nobre *et al*., 2020, 2022; Araujo *et al*., 2023]. Various methods can be employed for community detection in networks with the characteristics of our proposed models [Souravlas *et al*., 2021]. In this work, we choose to use the Louvain algorithm [Blondel *et al*., 2008], which has been widely adopted in the literature.

The Louvain algorithm aims to maximize modularity Q, a measure that quantifies the quality of a network partition into communities [Newman and Girvan, 2004]. Intuitively, modularity captures the internal density of nodes within each community compared to the organization of vertices in a random network with the same degree distribution. For a weighted graph $G = (V, E)$, where $w(e_{ij})$ represents the weight of the edge connecting vertex v_i to vertex v_j, and k_i represents the sum of the edge weights connected to v_i. Equation 7 presents the modularity function used by the Louvain algorithm to evaluate the quality of the network partition into communities.

$$Q = \frac{1}{2M} \sum_{v_i, v_j \in V} \left[w(e_{ij}) - \frac{k_i k_j}{2M} \right] \delta(c_i, c_j), \tag{7}$$

where M is the sum of all edge weights in G, c_i is the community assigned to vertex v_i, and δ is the Kronecker delta, a function that returns 1 if its operands are equal and 0 otherwise.

Modularity is constrained to the range $Q = [-1/2, +1]$. Values above 0.3 are strong evidence of well-structured communities [Newman and Girvan, 2004]. Finding the optimal community structure is an NP-hard problem, and heuristic algorithms, such as Louvain, are commonly used as alternatives. The Louvain algorithm follows an agglomerative strategy, initially assigning each of the n vertices in the network to a separate community. At each step, the algorithm reassigns vertices to communities to maximize modularity gain. Once no further gain is possible, each community is treated as a single node, and the process is repeated. The intermediate structure with the highest modularity is returned as the final output.

3.4.2 Topic Modeling

To identify and analyze the topics discussed in social interactions, we apply the widely used BERTopic [Grootendorst, 2022], combining deep learning, dimensionality reduction techniques, and density-based clustering to identify recurring themes in extensive textual data. Unlike traditional approaches, BERTopic captures textual semantics more effectively, providing a richer representation of emerging topics. In our approach, topic modeling was applied to the complete set of messages, without incorporating any segmentation based on edge classifications. The bidimensional structure was introduced only after the topics had been identified, enabling the analysis of how different dissemination profiles contribute to the spread of specific themes. This separation ensures that the topic modeling results remain unbiased by network structure and allows for an indirect comparison with traditional, unsegmented interaction networks.

The BERTopic process begins by converting texts into vector representations using the Sentence-BERT (SBERT) framework [Reimers, 2019]. Given that our case studies focus on Portuguese-language datasets, we leverage the BERTimbau model [Souza *et al*., 2020], a pre-trained BERT-based architecture fine-tuned to better capture semantic nuances in Portuguese [de Souza, 2020; Silva and Freitas, 2022]. These vector representations preserve semantic relationships between words and phrases, forming a robust foundation for subsequent analysis.

To prepare the input texts for embedding and topic modeling, we applied a uniform preprocessing pipeline to all messages, regardless of sentence count. Messages containing multiple sentences were treated as a single document and encoded as such by the Sentence-BERT model. No sentence segmentation was performed, as the transformer-based embedding framework is designed to capture contextual semantics across entire passages. Additionally, all texts were lowercased, emojis and special characters were removed, user mentions (e.g., @username) were filtered out, and extraneous whitespace was cleaned. These steps ensure that each message—whether single- or multi-sentence—is normalized and semantically represented in a consistent manner.

Once the vectors are generated, we apply Uniform Manifold Approximation and Projection (UMAP) [Armstrong *et al*., 2021] to reduce dimensionality, optimizing clustering without losing essential semantic information. This dimensionality reduction is crucial for improving clustering efficiency while preserving local and global relationships among texts. While we employ BERTimbau in this study to align with our Portuguese-language case studies, our framework is model-agnostic and can accommodate any language model suited to the dataset under analysis.

$$W_{x,C} = \|\mathrm{tf}_{x,C}\| \times \log\left(1 + \frac{A}{f_x}\right), \tag{8}$$

where:

- $W_{x,C}$ represents the relevance of word x in cluster C,
- $\|\mathrm{tf}_{x,C}\|$ is the frequency of word x within cluster C,
- f_x is the total frequency of word x across all clusters,
- A is the average number of words per cluster.

To enhance keyword diversity, BERTopic applies the Maximum Marginal Relevance (MMR) criterion [Gunawan *et al.*, 2019; Mao *et al.*, 2020] to select representative terms for each topic that are both relevant and minimally redundant. This strategy ensures that extracted keywords capture distinct aspects of the topic, rather than being overly similar to one another. We follow the tuning guidelines recommended in the official BERTopic documentation[2], adjusting parameters to produce a representative number of coherent topics without excessive fragmentation.

MMR achieves this by maximizing the relevance of a candidate keyword to the topic while simultaneously promoting diversity relative to the keywords already selected. Equation 9 presents the formulation of MMR as a function of the topic vector q:

$$\mathrm{MMR}(q) = \arg\max_{x \in S}\left[\lambda \cdot \mathrm{Sim}(x, q) - (1 - \lambda) \cdot \max_{y \in R} \mathrm{Sim}(x, y)\right] \tag{9}$$

where S is the set of candidate keywords, R is the set of already selected keywords, q represents the topic embedding, and $\lambda \in [0, 1]$ controls the trade-off between relevance and diversity. In BERTopic, the similarity functions $\mathrm{Sim}(x, q)$ and $\mathrm{Sim}(x, y)$ are computed using cosine similarity between the respective word embeddings. This formulation ensures that newly selected keywords are not only relevant to the topic but also semantically distinct from previously chosen terms.

After the initial topic inference, we apply two post-processing procedures to refine topic quality and reduce noise. First, we apply the `reduce_outliers` method to reassign documents initially marked as noise (e.g., by HDBSCAN) or weakly associated with any cluster. This is done using semantic similarity between document embeddings and topic centroids computed via c-TF-IDF, ensuring that even initially uncertain texts are incorporated meaningfully into the topic space. Second, we apply the `reduce_topics` method to merge semantically similar topics, especially when fragmentation leads to redundancy or thematic overlap across clusters. These steps reduce topic sparsity and facilitate downstream interpretability. Parameter values were selected based on BERTopic guidelines[3] and qualitative evaluation of topic coherence, aiming to balance interpretability and cluster granularity across datasets. The specific configurations used for each case study are explicitly reported in the corresponding sections.

To reduce computational cost, especially when working with millions of posts, we deduplicate messages prior to embedding. This preprocessing strategy avoids redundant encoding operations while retaining the semantic breadth of the data. After topic modeling, we reassign topic labels back to all duplicated messages using a one-to-one mapping with the original set. This maintains full corpus coverage in downstream analyses, such as topic dissemination mapping. Once topics are extracted, each input text (tweet or Telegram message in our case studies presented in the next section) is assigned a topic, allowing us to examine its dissemination patterns. We assess whether specific themes are amplified predominantly at high speed, high volume, or a combination of both. This enables us to understand which topics gain traction on social media, how they spread, and which dissemination patterns play key roles in this process.

4 Case Studies

This section describes two case studies explored using our framework, as outlined in the previous section.

4.1 Twitter/X

Our first case study analyzes information dissemination on Twitter, recently re-branded as X[4]. Twitter/X has been widely used for information dissemination, particularly in political events [Pacheco *et al.*, 2020; da Rosa *et al.*, 2022; Linhares *et al.*, 2022; Pohl *et al.*, 2023; Saxena *et al.*, 2023]. Here, we focus our analysis on events related to the 2022 Brazilian elections, using a dataset provided by G. Da Fonseca *et al.* [2024]. This dataset consists of data collected between October 1st, 2022, and February 10th, 2023, covering key events of the 2022 Brazilian general elections, such as the two voting rounds, vote counting, electoral polls, debates, and the attack on the Chamber of Deputies, Senate, and Supreme Federal Court buildings on January 8, 2023. During this period, the spread of anti-democratic content on Twitter was already observed[5].

To analyze user interactions, we consider network models for two daily time windows $T = \{01/11/2022, 08/01/2023\}$, corresponding to the two days with the highest observed activity. Thus, we have the graphs $G^{01/11/2022}_{\text{volume}}$ and $G^{01/11/2022}_{\text{time}}$, representing users active on November 1, 2022, and their respective co-interactions on that date, as well as the graphs $G^{08/01/2023}_{\text{volume}}$ and $G^{08/01/2023}_{\text{time}}$, corresponding to users active on January 8, 2023, and their respective co-interactions in that time window. The edge weights $w(e_{ij})$ between users v_i and v_j are computed according to the model defined in Section 3.1.

The parameters for the *Polya Urn* method in backbone extraction for this case study were set according to the guidelines of [Marcaccioli and Livan, 2019; Gomes Ferreira *et al.*, 2022], using a reinforcement parameter of $a = 0.5$ for both

[2] https://maartengr.github.io/BERTopic/getting_started/parameter%20tuning/parametertuning.html

[3] https://maartengr.github.io/BERTopic/getting_started/parameter%20tuning/parametertuning.html

[4] www.x.com

[5] https://portal.stf.jus.br/noticias/verNoticiaDetalhe.asp?idConteudo=508935&tip=UN,https://www.bbc.com/portuguese/brasil-63990040

G_{volume} and G_{time}, with $\alpha = 0.05$ for G_{volume} and $\alpha = 0.1$ for G_{time}.

Regarding the Twitter/X case study reproducibility, we strictly followed the platform's developer policy, which only permits the redistribution of tweet identifiers (tweet IDs), not full tweet content. Consequently, the dataset used in our analyses is shared in the form of tweet IDs, enabling other researchers to reconstruct the full messages via the official API, in accordance with Twitter's terms of service [6].

4.2 Telegram

Our second case study focuses on the dissemination of information in Telegram groups. The platform connects users through end-to-end conversations and groups, proven effective in large-scale information dissemination [Dargahi Nobari *et al.*, 2017; Urman and Katz, 2020; Castagna *et al.*, 2023; Ulizko *et al.*, 2022; Ng *et al.*, 2024]. In this case study, we use messages from politically oriented Telegram groups in Brazil, posted between September 1st, 2022, and January 31st, 2023, covering the period of the presidential elections and the January 2023 riots. This dataset was provided by [Venâncio *et al.*, 2024a].

To analyze how users disseminate content in Telegram groups, we use the network models initially applied by the authors, which connect users who share the same textual content across observed groups. Additionally, we analyze three 15-day time windows, considering seven days before and seven days after the following key events: $T = \{02/10/2022\ (1^{\text{st}}\ \text{Round}), 30/10/2022\ (2^{\text{nd}}\ \text{Round}), 08/01/2023\ (\text{Riots})\}$, representing the two election rounds and the riots of January 8, 2023.

We define the network models for these time windows as follows: $G^{1^{\text{st}}\ \text{Round}}_{\text{volume}}$ and $G^{1^{\text{st}}\ \text{Round}}_{\text{time}}$, $G^{2^{\text{nd}}\ \text{Round}}_{\text{volume}}$ and $G^{2^{\text{nd}}\ \text{Round}}_{\text{time}}$, and $G^{\text{Riots}}_{\text{volume}}$ and $G^{\text{Riots}}_{\text{time}}$, corresponding to users who posted at least one message in a Telegram group during the specified time window. Each edge e_{ij} indicates that users v_i and v_j shared the same message (identical textual content) at least once within the same time window. The parameters for the *Polya Urn* method in backbone extraction were again set according to the guidelines of [Marcaccioli and Livan, 2019; Gomes Ferreira *et al.*, 2022], using a reinforcement parameter of $a = 0.5$ and $\alpha = 0.05$ for G_{volume} and $\alpha = 0.1$ for G_{time}.

The Telegram case presents more delicate ethical and legal challenges. While all data were collected exclusively from public groups, Telegram's platform policy does not explicitly define whether the redistribution of user-shared messages is permitted, especially when such content may contain personal identifiers such as names, phone numbers, or references to organizations. Given the politically sensitive context of the events analyzed, such as the 2022 Brazilian elections and the January 8, 2023 riots, we opted not to release the raw textual content of the messages. Nevertheless, we have made the structural components of our analysis, including the network models, the backbone structures and the edge classification code, publicly available in the repository mentioned above. These materials allow for full reproduction of the structural and topological analyses. Once again, the only component that remains unreproducible without access to the raw messages is the textual analysis phase, which necessarily depends on message content. All data were accessed via the Telethon API[7], strictly within public groups and without requiring authentication or user credentials. This approach ensures that the data collection complies with the platform's public access model, while respecting ethical boundaries associated with user privacy and content sensitivity.

[6] Twitter Developer Policy: https://developer.x.com/en/developer-terms/policy

Table 1. Structural characteristics of the Twitter/X and Telegram networks.

Network	Date	n	m	d	$\hat{k}$	ACC	CC	# $Com.$	Q
Twitter/X	08/01/2023	28,247	86,478,838	0.216	6123	0.689	1	5	0.23
Twitter/X	01/11/2022	27,920	72,777,139	0.186	5213	0.700	1	5	0.49
Telegram	1st Round	1,625	81,236	0.062	100	0.641	80	110	0.12
Telegram	2nd Round	3,444	161,383	0.027	94	0.621	164	193	0.12
Telegram	Riots	4,772	71,046	0.006	30	0.549	135	155	0.35

5 Results

This section presents the results of applying the methodology detailed in Section 3 to the case studies described in Section 4, considering the proposed framework for bidimensional analysis of information dissemination on social media platforms in real-world contexts. After a topological characterization of the studied networks, we present the analysis of dissemination patterns from both a structural and a content-based perspective.

5.1 Topological Analysis of the Networks

Table 1 summarizes key topological properties of the dissemination networks for each case study. The number of nodes (n) corresponds to unique users, and the number of edges (m) reflects co-interactions based on retweets (Twitter/X) or posting identical messages (Telegram). The average degree ($\hat{k}$) indicates the mean number of connections per user, while the clustering coefficient (ACC) measures the tendency of nodes to form tightly knit groups. Network density (d) is the ratio of observed to possible edges. CC refers to the number of connected components in the network, while # $Com.$ and Q denote the number of communities and modularity, respectively, used to evaluate community structure.

By comparing the Twitter/X and Telegram networks, we observe distinct structural characteristics that reflect differences in user behavior and platform affordances. In the Twitter/X networks, edges between nodes are formed through shared retweets, a behavior that occurs at high frequency, resulting in densely connected networks with high average degree. In contrast, Telegram lacks a native repost function as edges are created only when users post nearly identical messages. Even so, the resulting networks display significant structure and community formation, albeit with lower density and smaller average degrees.

The high density observed in both case studies (especially in Twitter/X) strongly indicates the need for filtering strategies to highlight the most relevant relationships to better understand the information dissemination phenomenon. This is achieved through the application of backbone extraction

[7] `https://docs.telethon.dev/en/stable/`

Table 2. Topological properties of Twitter/X networks per dissemination class.

Class	Date	n	m	d	$\hat{k}$	ACC	CC	# *Com.*	Q
1	01/11/2022	27,911	47,173,429	0.121	3380	0.467	1	6	0.49
2	01/11/2022	20,363	1,373,833	0.007	135	0.586	17	24	0.43
3	01/11/2022	27,899	23,980,819	0.061	1719	0.356	2	4	0.47
4	01/11/2022	18,453	249,058	0.001	27	0.271	29	48	0.62
1	08/01/2023	28,230	41,729,887	0.104	2956	0.300	1	5	0.26
2	08/01/2023	27,401	4,940,636	0.013	361	0.397	12	21	0.31
3	08/01/2023	28,236	37,192,895	0.093	2634	0.389	1	6	0.32
4	08/01/2023	27,578	2,615,420	0.007	190	0.350	8	14	0.34

in this study. The presence of a large number of edges, potentially including a significant proportion of noisy relationships, also affects the identification of community structures within the networks. This is reflected in the modularity values, which indicate weakly defined groups. Even in Telegram networks, which exhibit highly fragmented components (as indicated by the large number of communities and connected components), we note relatively low modularity values, corresponding to loosely structured partitions, as in the case of the Riots time window.

5.2 RQ1: Classification and Characterization of Dissemination Edges Profiles

Based on the bidimensional classification of dissemination, considering both volume and interaction speed, a topological analysis is performed to understand the characterization of each profile within the studied contexts. This analysis is enabled after applying the initial steps of the proposed *framework*, as described in Section 3. Table 2 presents the topological characteristics of the networks, as well as the community structures for each class obtained after identifying the *backbones* in both dimensions (volume and speed) for Twitter/X.

After identifying the classes for the volume and speed dimensions, the network models retain the same set of vertices, with potential variations in the set of edges representing different dissemination profiles. Thus, although there are differences in the number of vertices across networks of different classes caused by the omission of vertices without interactions in specific profiles, a significant balance in the representation of users in each class can be observed. On the other hand, there is a clear imbalance in the number of edges, reflecting a considerably higher prevalence of interactions in `Class 1`. A clearer understanding of the characteristics of the dissemination profiles identified for each class can be obtained by observing Figure 2, which illustrates the distribution of vertices and edges in the networks representing the interactions of each class for Twitter/X.

While the network for `Class 1`, which represents relationships that are not salient in either the volume or speed dimensions, exhibits a high number of edges, the network for `Class 4`, which represents relationships that are exceptionally relevant in both dimensions, presents a significantly lower number of edges, especially in the time window corresponding to 01/11/2022. This is reflected in some of the statistics shown in Table 2, such as the average degrees (and consequently, the densities), which are substantially higher for `Class 1` and lower for `Class 4`, once again highlighting the importance of identifying the most relevant edges for understanding the phenomenon of information dissemination. It is also interesting to note a difference between the characteristics of the networks of `Class 2`, which preserves salient edges related to volume, and `Class 3`, which preserves salient edges related to speed. These results indicate that the volume dimension presents much more dispersed relationships than the speed dimension, making the networks of `Class 2` substantially less dense than those of `Class 3`.

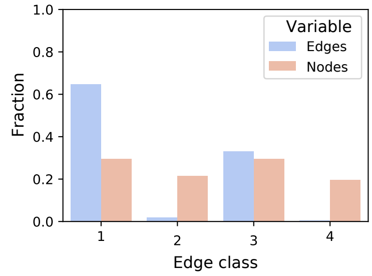

(a) Time window: 01/11/2022

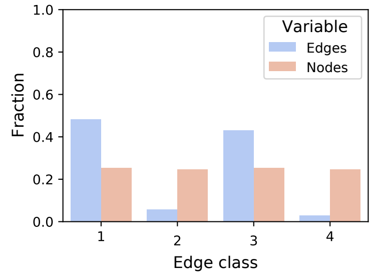

(b) Time window: 08/01/2023

Figure 2. Fraction of nodes and edges per dissemination class in Twitter/X networks.

The organization of vertices into communities in networks with different edge classes further reinforces the importance of applying the *backbone* in a bidimensional manner. This becomes evident when observing the modularity values of `Class 4` compared to other classes in both time windows, particularly in the one related to 01/11/2022. When noisy edges, potentially representing sporadic or random relationships, are filtered out while simultaneously considering both the volume and speed dimensions, a clearer user community structure emerges.

To provide a deeper visualization of how information propagates in networks with edges from different classes, Figures 3 present the cumulative distribution functions (CDFs) of the number of *retweets* and the average *retweet* time for each class. The x-axis represents, in logarithmic scale, the values of the random variable in question (either the number of *retweets* or the average time). At the same time, the y-axis indicates the cumulative probability of that variable.

Figures 3(a) and 3(b), which depict the distribution curves for the volume of *retweets* per edge class, show that `Classes 2` and 4 exhibit distributions significantly shifted to the right on the x-axis. This indicates that the most relevant interactions in terms of volume, identified through the *backbone* method, also correspond to a higher number of *retweets*. In Figure 3(a), for example, approximately 65% (80%) of the edges in `Class 1` (`Class 3`) have a weight of 1. Conversely, for `Classes 2` and 4, which capture edges with weights representing a larger number of co-interactions, the minimum observed values start at 3.

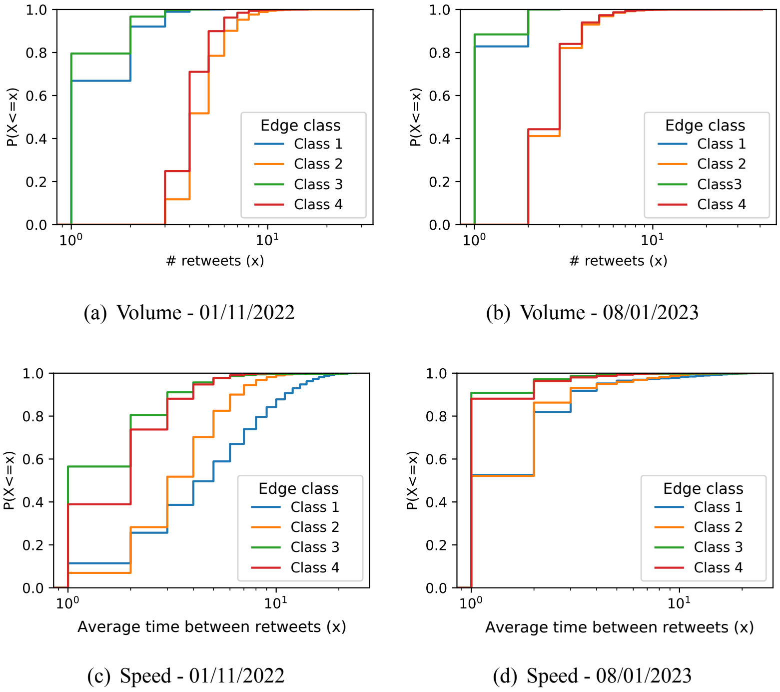

(a) Volume - 01/11/2022 (b) Volume - 08/01/2023

(c) Speed - 01/11/2022 (d) Speed - 08/01/2023

Figure 3. CDF of *retweets* and average time between *retweets* for each class in the time windows 01/11/2022 and 08/01/2023.

An analogous analysis can be conducted based on the time dimension, as shown in Figures 3(c) and 3(d), which display the CDFs of the average time between *retweets* in minutes for each class. To enhance the intuitive understanding of the analysis, the time normalization defined in Equation 2 (Section 3) was disregarded in this experiment. `Classes 3` and 4 exhibit significantly lower average *retweet* times, demonstrating that the most relevant edges in terms of the time dimension, captured within these classes, represent faster dissemination, regardless of whether it occurs at low or high volume. For instance, in Figure 3(d), approximately 90% of the edges in `Classes 3` and 4 were generated with an average co-interaction time, i.e., co-*retweets*, of around 1 minute[8]. This suggests a much higher probability of extremely rapid dissemination compared to `Classes 1` and 2, where this figure is approximately 52%.

Considering the dissemination of messages occurring on Telegram, Table 3 presents the topological characteristics of the networks, as well as the community structures for each class obtained after the identification of the *backbones* in both dimensions (volume and speed).

Unlike the networks observed in the Twitter/X case study, the Telegram networks exhibit not only an imbalance in the number of edges but also a significant discrepancy in the number of vertices among the classes. Figure 4, which graphically illustrates the distribution of vertices and edges among classes, helps us better understand the topological characteristics of these networks.

In Telegram networks, `Class 1`, which corresponds to edges not filtered by the *backbone* method in either dimension, concentrates a considerably larger number of nodes and edges, similar to what was observed in the Twitter/X case study. However, other topological characteristics are not repeated in this case study. `Class 2` and `Class 4`, representing salient relationships in terms of volume, exhibit a notably low number of edges and vertices, revealing some interesting aspects of the information dissemination phenomenon in this social network.

Table 3. Topology of Telegram networks by dissemination class.

Class	Date	n	m	d	$\hat{k}$	ACC	CC	$\#Com.$	Q
1	1st Round	1,406	54,470	0.055	77	0.449	118	137	0.13
2	1st Round	53	35	0.025	1	0.056	19	20	0.88
3	1st Round	1,208	26,446	0.036	44	0.419	2	7	0.46
4	1st Round	120	288	0.040	5	0.205	9	15	0.38
1	2nd Round	3,150	107,218	0.022	68	0.488	218	237	0.16
2	2nd Round	93	78	0.018	2	0.063	21	24	0.84
3	2nd Round	2,141	53,743	0.023	50	0.443	3	6	0.46
4	2nd Round	137	344	0.037	5	0.208	5	11	0.41
1	Riots	4,763	70,313	0.006	30	0.528	136	158	0.35
2	Riots	254	427	0.013	3	0.173	23	33	0.55
3	Riots	263	303	0.009	2	0.092	19	26	0.79
4	Riots	5	3	0.300	1	0.000	2	2	0.44

Firstly, the low number of edges that represent the sharing of a large volume of the same content shows that this type of behavior is highly concentrated among a few pairs of users. This fact is strongly related not only to the information dissemination phenomenon but also to how the network was modeled on Telegram, where the sharing of messages with identical textual content defines an edge. For example, this action is less intuitive than a *retweet* on Twitter/X. Furthermore, the results for `Class 2` and `Class 4` also indicate that the vertices associated with these classes concentrate relationships only within this specific characteristic and, in this sense, are less diverse than those in Twitter/X. As a result, the networks in these classes exhibit high fragmentation besides a low number of vertices and edges, which can be observed in the high number of connected components shown in Table 3.

[8]The Ceiling function was used to improve visualization and prevent the average time from being zero.

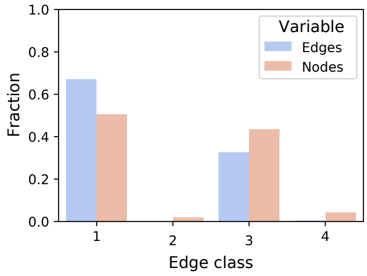

(a) Time window: 1st Round

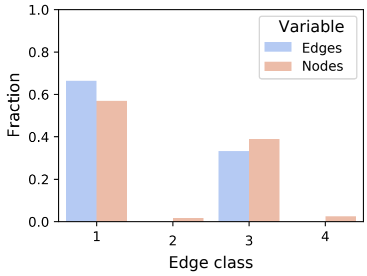

(b) Time window: 2nd Round

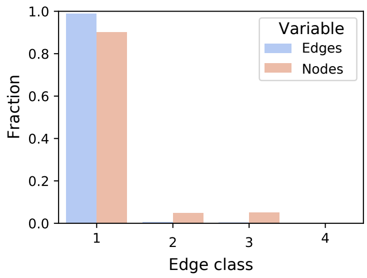

(c) Time window: Riots

Figure 4. Fraction of nodes and edges per dissemination class on Telegram.

The fragmentation observed in the networks of different classes on Telegram directly impacts the identified community structure. In many cases, modularity values above 0.8 are observed, such as in `Class 2` during the time windows of the 1st Round and 2nd Round. However, this result should be considered with caution when analyzing the phenomenon of information dissemination, as the clear definition of the community structure may result from the division of vertices into distinct components. Thus, even if information spreads within communities, each component will potentially contain a small number of vertices, making message dissemination more difficult, a limitation further reinforced by the barriers imposed by different components.

This finding indicates that considering edges in the *backbone* identified only in the volume dimension may not be an appropriate strategy for understanding the phenomenon of information spread on Telegram. On the other hand, the networks in `Class 3`, which retain edges identified in the *backbone* based on the time dimension, appear to exhibit better-organized communities, even with lower modularity values. These networks are notably less fragmented, as evidenced by their number of connected components. They can be structured into communities that enable a more in-depth investigation of the dissemination process. When comparing the results observed for Telegram with those found for Twitter/X, we gain an understanding of the complexity of information dissemination across different dimensions and platform-specific characteristics. This finding further reinforces the central argument of this study, emphasizing the importance of analyzing relevant relationships in social networks while considering not just one but multiple dimensions simultaneously.

To investigate how propagation occurs across different edge classes, Figure 3 presents the cumulative distribution functions (CDFs) of the number of messages and the average sharing time (in minutes) for each class and time window. Analyzing the volume of information through the distribution of the number of messages (Figures 5(a) and 5(b)), it is possible to observe that `Classes 2` and 4 exhibit significantly higher values, indicating edges with a high volume of dissemination. This suggests that the most relevant edges, present in the *backbone*, also represent relationships with high co-interaction. On the other hand, the edges in `Classes 1` and `3` show a significantly lower dissemination volume, with approximately 95% of edges weighting 1 across all scenarios. This indicates the strong presence of sporadic but rapid co-interactions.

Regarding the speed of information propagation on the platform, the CDFs in Figures 5(d), 5(e), and 5(f) reveal propagation patterns distinct from those observed on Twitter/X. Edges in `Classes 3` and 4 again exhibit lower average sharing times compared to other classes, especially in Figures 5(d) and 5(e). However, certain values in `Class 2` stand out, showing average times as low as those in `Classes 3` and 4. Recall that the *Polya Urn backbone* extraction method (Section 3) is based on a model that assumes that salient edges of a node are those whose weights significantly deviate from others in relation to the edges incident on the same node, meaning it follows a local analysis approach. This result demonstrates that, in the Telegram dissemination network, some edges in `Class 2` have weights as low as those in `Classes` 3 and 4, yet they are not captured by the *backbone* and, therefore, are not classified into these dissemination profiles. This occurs because these edges belong to pairs of nodes whose exclusive sets of edges exhibit distributions of average sharing times that are predominantly lower compared to the same distribution across the entire network. In these cases, the even lower values of these edge weight distributions cause some edges to be classified as belonging to `Classes 3` or 4, while others—slightly higher in relation to the node but significantly lower relative to the overall network pattern—are classified as belonging to `Class 2`.

In the Twitter case study, while Classes 2 and 4 reflect intense interactions aligned with traditional backbone filtering strategies with one dimension, we find that Classes 1 and 3 are crucial to understanding the distribution of narratives. For instance, Class 3 captures reactive amplification of condemnatory or news-oriented content, whereas Class 1 includes more persistent discussions around military intervention or ideological framing, often absent from high-volume dissemination networks. This suggests that traditional single-backbone approaches—typically recovering edges in Classes 2 and 4—may overlook important dissemination dynamics associated with slower or lower-volume interactions.

In summary, these node pairs have sets of edges that ex-

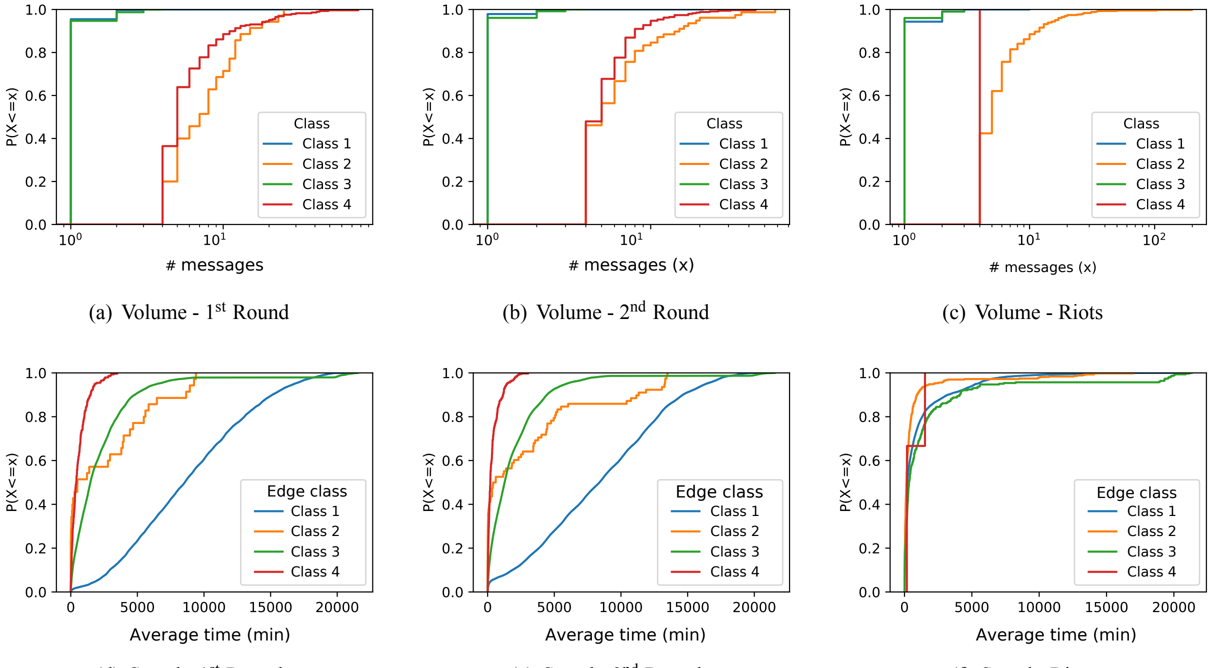

(a) Volume - 1st Round (b) Volume - 2nd Round (c) Volume - Riots

(d) Speed - 1st Round (e) Speed - 2nd Round (f) Speed - Riots

Figure 5. Cumulative distribution function of the number of messages and average sharing time for each class in Telegram networks.

hibit a high dissemination potential in terms of speed compared to other nodes in the network. Even the edges not captured by `Classes 3` and 4 still exhibit lower dissemination times compared to the rest. Consequently, our results reveal the existence of nodes in the network with predominantly faster dissemination values than those observed in the network as a whole. Finally, it is noteworthy that, during the riots, the distribution of weights based on average sharing time (Figure 5(f)) shows generally lower times for all classes.

5.3 RQ2: Topic Dissemination Driven by Edges' Classes

Following the methodological steps outlined in Section 3, the proposed framework enables the characterization of the topological structure of information dissemination networks on social media platforms and the analysis of the content shared within these environments. To maintain analytical coherence and avoid an excessive number of comparisons across multiple time windows, we focus our analysis on specific temporal snapshots for each case study. For Twitter/X, we examine data from January 8, 2023, a date marked by coordinated information flows surrounding the riots [Malagoli *et al.*, 2024]. For Telegram, we select the second round of the 2022 Brazilian presidential election, given its relevance in shaping political discourse [Venâncio *et al.*, 2024b,a].

In this analysis stage, we employ the BERTopic framework to identify and categorize topics based on textual patterns extracted from user interactions. For both platforms, we utilize the BERTimbau model[9] [Souza *et al.*, 2020], a pre-trained transformer-based language model optimized for Portuguese, ensuring high-quality semantic representations of textual data. The topic modeling process follows the dimensionality reduction step using UMAP, configured to balance local and global semantic structures (Section 3). For Twitter/X, we set $n_{neighbors} = 3$ and $n_{components} = 10$ to capture fine-grained topic distinctions, while for Telegram, a broader semantic grouping is applied with $n_{neighbors} = 5$. Clusters are then identified using HDBSCAN, with a minimum cluster size of 25 for Telegram, to ensure robust topic formation. To refine term selection, we apply TF-IDF vectorization with $min\_df = 0.01$ (removing terms appearing in fewer than 1% of documents) and $max\_df = 0.99$ (excluding terms present in more than 99% of documents) for both scenarios.

The topics identified in Table 4 reveal a range of narratives surrounding the events of January 8, 2023, on Twitter/X. Some topics reflect support for the demonstrations and their participants, such as *Topic 6*, which includes messages suggesting the presence of infiltrators as a means to shift blame away from protesters, forming a potential defense narrative. Similarly, *Topic 1* contains references to calls for military intervention in Brasilia to address the unfolding events. However, within the same topic, the presence of the term *coup* suggests that the discourse was not unidirectional, incorporating messages that framed the events as an attempted coup.

Given the nature of Twitter/X during data collection, with a substantial presence of journalists, media professionals, and news outlets, many messages appear to have an informative intent. This is particularly evident in *Topics 3, 7, and 11*, which pertain to reports on the invasions, mentioning terms such as *Twitter*, *videos*, and *news*, as well as direct references to the locations targeted during the event, including *Planalto*, *Palace*, and *STF*. Despite their apparent informational nature, these messages also include terms indicative of a condemnatory stance, such as *invasion*, *coup plotters*, and explicit ref-

[9] https://huggingface.co/neuralmind/bert-large-portuguese-cased

ID	Key Words	Description
1	Brasilia, intervention, military, coup, democracy, congress	Calls for military intervention and allegations of a coup in Brasilia
2	Governor, Ibaneis, security, federal, invasion, government	Discussion on the Governor of the Federal District's reaction to the events
3	Sunday, invasion, STF, Planalto, Brasilia, Three Powers, president	General discussion on the events of the day
4	Amnesty, Brasilia, prisoners, businessmen, funding, authorization	Debate on the possible funding of protesters and calls for arrests without amnesty
5	People, illegitimate, coup plotters, anti-democratic, criminals, respect	Accusations that the events were illegitimate and anti-democratic
6	Infiltrators, riot, strange, break, attack, lie	Allegations of infiltrators creating chaos to blame protesters
7	Congress, coup plotters, palace, group, politics, government, STF, Bolsonaro, military	Reports specifying locations attacked during the event
8	Dino, PRF, PF, GDF, contain, together, justice, act, defense	Calls for action by federal and state security forces
9	Special forces, authorize, use, ministries, government, attitude	Requests for federal special forces to intervene
10	Fans, Gaviões, Galoucura, go, Brasilia, defend, want	Mentions of organized football fan groups as an alternative to contain protests
11	Twitter, videos, Brasilia, defense, about, news, show	Sharing news and videos about the events on social media
12	Operation, defend, start, war, against, urgent	Associating the events with warfare and emergency actions
13	Vandalism, right-wing, dictatorship, fascism, heritage, crime	Associating property destruction with political extremism
14	President, France, support, White House, Biden, anti-democratic	Mentions of international leaders and foreign perspectives
15	People, serious, population, country, life, all, impacted	Expressions of concern about the severity and impact of the events

Table 4. Discussion topics found on Twitter regarding the events of January 8.

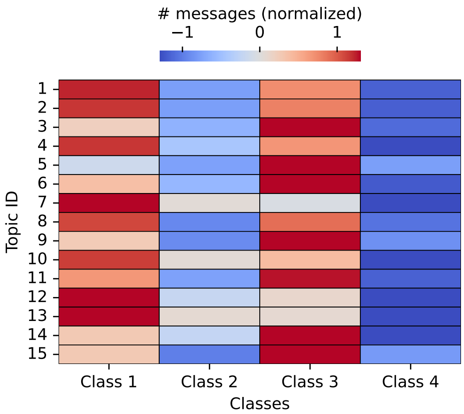

Figure 6. Heatmap of topic distribution across classes on Twitter.

erences to former president *Bolsonaro*, suggesting a broader trend of denunciation in the discourse.

A more explicit tendency toward condemnation is observed in the majority of the identified topics. *Topic 5*, for instance, is characterized by strong terms such as *illegitimate*, *coup plotters*, *anti-democratic*, and *criminals*, indicating a narrative that frames the events as an attack on democratic institutions. Similarly, *Topic 13* establishes a direct association between the events and right-wing extremism, as evidenced by terms such as *vandalism*, *right-wing*, *fascism*, and *crime*.

Another distinctive element is the presence of references to football fan groups, such as *Galoucura* and *Gaviões* (*Topic 10*). This appears to allude to the role of these groups in countering road blockades set up by demonstrators dissatisfied with the election results, a phenomenon widely documented in Brazil's political landscape. Given the often ironic and satirical nature of discourse on Twitter/X, this topic may reflect sarcastic remarks or discussions about the role of these groups in the broader political context. Taken together, these findings indicate that while supportive and defensive narratives existed, the predominant discourse on Twitter/X regarding the events of January 8 was one of denunciation and criticism, with a strong emphasis on characterizing the acts as illegitimate, anti-democratic, and criminal.

To analyze how different dissemination profiles contribute to topic amplification on Twitter/X, we normalize the message distribution within each topic using a *z-score* transformation (Figure 6)[10]. This row-wise normalization emphasizes the relative prominence of each dissemination profile within individual topics, enabling direct comparison between classes while controlling for absolute message volume. Positive values indicate class overrepresentation within the topic, whereas negative values denote underrepresentation.

Examining Figure 6, we observe distinct patterns in how topics are distributed across dissemination profiles. In *Class 1* (low-volume, low-speed dissemination), the strongest presence is observed in *Topics 1, 2, 4, 7, 8, 10, 12, and 13*. These topics primarily involve discussions on military intervention, security measures, and the broader context of the January 8 events. The prominence of *Topic 1* (military intervention) and *Topic 2* (government responses) suggests that these discussions were widely spread within low-engagement groups, while *Topic 10* (football fan groups as alternative security forces) and *Topic 12* (war-like rhetoric) indicate that this group engaged in more general discussions about the nature of the events.

In *Class 2* (high-volume, low-speed dissemination), there is a notable presence of *Topics 7, 10, and 13*, albeit with less intensity than in Class 1. *Topic 7* (discussion of key locations and political figures) suggests that this group maintained a sustained discussion about the events, while *Topic 13* (associating the events with political extremism) reflects how high-volume disseminators played a role in reinforcing ideological interpretations over time.

A different pattern emerges in *Class 3* (low-volume, high-speed dissemination), where *Topics 3, 5, 6, 9, 11, 14, and 15* are more prominent. This suggests that users in this category were engaged in rapidly amplifying discussions centered on general event reports (*Topic 3*), condemnation of the acts (*Topic 5*), accusations of infiltration (*Topic 6*), calls for security forces to act (*Topic 9*), and news dissemination (*Topic 11*). The presence of *Topics 14 and 15*, which reference international perspectives and broader concerns, further indicates that this dissemination profile was associated with a more reactive engagement with the unfolding events. Finally, *Class 4* (high-volume, high-speed dissemination) does not exhibit a strong presence in any particular topic. The rel-

[10]Standardized as: $Z_j = \frac{X_j - \mu}{\sigma}$, where X_j is the message count for class j within a given topic, μ is the mean count across classes for that topic, and σ is the corresponding standard deviation.

atively balanced distribution across topics suggests that mass disseminators on Twitter/X engaged with a variety of discussions but did not play a disproportionate role in amplifying any specific narrative.

These findings suggest that different dissemination profiles played distinct roles in shaping online discourse. While low-speed, high-volume disseminators contributed to sustaining broad discussions about security and ideological narratives, rapid amplifiers were particularly engaged in reactive and event-driven discourse. Meanwhile, high-speed, high-volume actors participated in a more evenly distributed manner, reinforcing multiple discussions without strong thematic preference.

Like the approach taken for Twitter/X, Table 5 presents the leading sets of keywords identified in Telegram discussions related to the 2022 Brazilian elections. These topics were extracted without yet considering dissemination profiles, providing an overall view of the themes circulating within the platform. Unlike Twitter/X, where denunciatory language predominated, Telegram exhibits a more balanced distribution between supportive, critical, and ostensibly informational messages. However, the presence of organizational and mobilization elements within certain topics suggests that Telegram was used not only as a discussion space but also as a platform for strategic coordination. *Topic 1*, for instance, contains explicit references to electoral fraud and government legitimacy, incorporating terms such as *fraud*, *patriots*, and *revolution*, which align with narratives questioning the legitimacy of the election results and the administration in power. This pattern is further reinforced in *Topic 5*, which associates military and religious rhetoric with calls for action, as evidenced by terms such as *armed*, *forces*, *SOS*, and *freedom*.

Notably, several topics suggest that Telegram messages were not solely aimed at reporting or commenting on the events but also served a strategic role in mobilization and coordination. *Topics 3, 4, and 9* indicate an effort to expand the dissemination of content across multiple digital platforms, as observed in the presence of terms such as *Bitchute*, *Gettr*, and *TikTok*, which are alternative media channels often used to circumvent content moderation policies on mainstream platforms. Additionally, *Topic 3* explicitly encourages individuals to join discussion groups on *WhatsApp* and *Telegram*, reinforcing the intent to build and maintain a network of aligned participants.

Beyond online mobilization, certain topics also reference offline actions. *Topic 12*, for instance, contains references to military terminology and geographical markers, such as *battalion*, *infantry*, and *avenue*, which suggest discussions related to physical presence and organization. Similarly, *Topic 2* includes terms such as *protest*, *alert*, and *streets*, which indicate direct calls for civil mobilization.

Although most Telegram messages align with pro-demonstration narratives, topics expressing criticism or opposition to the events can also be identified. In some cases, this criticism appears in a mocking or derogatory tone, as exemplified by *Topic 8*, which contains terms such as *cattle*, *cry*, and *embarrassment*, reflecting pejorative language used by opponents of former president Jair Bolsonaro to refer to his supporters. Other topics adopt a more assertive critical stance, as seen in *Topic 6*, which raises concerns about censorship and impeachment, and in discussions containing terms such as *decree*, *shame*, and *Alexandre de Moraes*, referring to the then-president of the Superior Electoral Court.

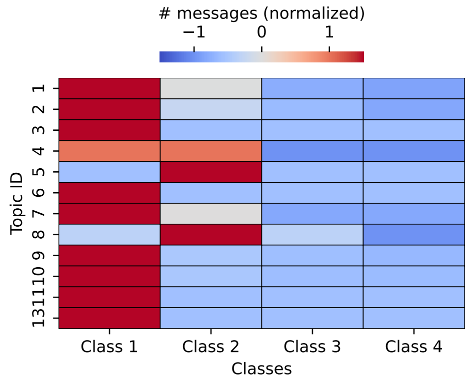

Figure 7. Heatmap of topic distribution across classes on Telegram.

Overall, the analysis of Telegram discussions during this period suggests that the platform was not only a space for discourse but also a tool for mobilization, with prominent narratives advocating for intervention, questioning electoral legitimacy, and coordinating actions both online and offline. Unlike Twitter/X, where condemnation and denunciation were dominant, Telegram's discourse landscape reflected a stronger presence of organized support for the events, intertwined with calls to action and alternative information dissemination strategies.

To analyze how different dissemination profiles contributed to Telegram topic amplification, we examine the distribution of messages across dissemination classes, as shown in Figure 7. To enable comparison of the relative prominence of each class within individual topics, we again applied row-wise *z-score* normalization. Remember that this transformation standardizes the values of each row (i.e., each topic) to have zero mean and unit standard deviation, thereby highlighting deviations relative to the topic's internal distribution. By controlling for absolute message volume, this approach emphasizes distinctive dissemination patterns within topics and facilitates their comparison across the dataset.

Examining Figure 7, we observe a clear predominance of Class 1 (low-volume, low-speed dissemination), which exhibits the highest z-score values across nearly all topics. This suggests that most Telegram messages were shared within communities characterized by slower, lower-intensity propagation, indicative of internally reinforced discourse rather than broad viral spread. For instance, Topics 1, 2, and 6 stand out with the strongest representation in Class 1. These topics include narratives around electoral fraud and government legitimacy (Topic 1), calls for mobilization (Topic 2), religious and criticism of censorship and impeachment rhetoric (Topic 6). This dissemination pattern suggests that Telegram was used to circulate ideologically charged narratives in a more insular and sustained manner, possibly reflecting strategic or grassroots coordination efforts.

Class 2 (high-volume, low-speed dissemination) shows relevance in a smaller set of topics, namely, Topics 4, 5, and 8. Topic 4 relates to the promotion of alternative media platforms (e.g., YouTube, Telegram, Bitchute), Topic 5 reinforces previously mentioned mobilization themes, and

ID	Key Words	Description
1	country, make, population, government, elections, will, ballots, Moraes, fraud	Discussion on electoral fraud and government legitimacy
2	share, patriots, attention, let's, day, streets, urgent, people, protest, alert, invade, front	Calls for street protests and civil mobilization
3	news, terrabrasilnoticias, elections, vistapatria, politics, aliadosbrasiloficial, gettr, telegram, uol, whatsapp, participate	Sharing news and political content through alternative media platforms
4	youtube, channel, telegram, watch, youtube, proxy, twitter, live, network, ballots, links, gettr, app, bitchute	Promotion of live streams and alternative social media networks
5	God, forward, leader, we can, together, forces, SOS, moment, let's, save, right, streets, armed, possible, guidance	Religious and military appeals for action and guidance
6	jornaldacidadeonline, news, see, video, live, survey, zambelli, censorship, pan, debate, moraes, carla, young, serious, impeachment	Criticism of censorship and calls for impeachment of authorities
7	status, twitter, brazil, live, biden, dictatorship, pix	Mentions of international politics and allegations of dictatorship
8	cattle, cry, embarrassment, bolsonarists, sad, accept, lost, regrettable, suffering	Sarcastic and critical remarks towards Bolsonaro supporters
9	youtube, new, speech, watch, official, channel, president, video, press conference, interview, dictatorship, stay, love, hands	Announcements and live streams of presidential statements
10	brazil, stolen, twitch, tv, video, gettr, material, signal, fraud, zambelli, crime, deception	Claims of election fraud and political deception
11	jungle, attack, northeast, give, chaos, lose, PT, audio, propaganda, second round, risk, Argentina, communism	Concerns about leftist policies and political instability
12	battalion, avenue, infantry, RS, SP, center, square, command, Caxias, Duque, arrested, military, army, street, captain	References to military actions and organization
13	videos, facebook, status, share, united, anomalies, twitch, instagram, save, report	Calls to share and document election-related anomalies

Table 5. Discussion topics found in Telegram groups regarding the 2022 Brazilian elections.

Topic 8 adopts a mocking tone toward political opponents, such as supporters of former president Bolsonaro. The presence of these topics in Class 2 indicates that while some content achieved higher overall volume, it was still disseminated gradually over time, suggesting sustained but non-viral sharing dynamics. In contrast, Classes 3 and 4 (high-speed dissemination, with low and high volume respectively) do not exhibit strong representation in any particular topic.

Compared to Twitter/X, where multiple dissemination profiles were active and distinct patterns emerged across all four classes—including rapid amplification of news and condemnatory content (Class 3) and broad, high-speed engagement (Class 4)—Telegram presents a more concentrated pattern. Here, slower, lower-volume dissemination (Class 1) dominates, and rapid or viral amplification is notably absent. While Classes 2 and 4—typically emphasized in traditional, volume-based backbone extraction methods—also appear in the Telegram data, they play a more limited role in shaping the thematic structure of the discussions. This highlights the value of incorporating both volume and speed dimensions, a one-dimensional approach focused solely on high-volume interactions would largely overlook the most active segments of the Telegram network. These results reinforce our core claim that different platforms afford different dissemination strategies, with Telegram functioning more as a space for ideological reinforcement and strategic mobilization, while Twitter/X supports broader, faster-paced information diffusion.

6 Conclusion

Social media platforms are now central to how modern society shares information, influencing everything from political debate to public mobilization. While their capacity for rapid, large-scale communication has reshaped social dynamics, it has also created significant challenges, including the spread of disinformation and the manipulation of public opinion. To address these issues, this paper introduced a bidimensional framework that moves beyond traditional, volume-only analysis by integrating both the volume and the speed of interactions. By applying this methodology to a case study on Telegram, we have demonstrated its effectiveness in identifying unique dissemination patterns and their role in content propagation.

Our findings confirm that different dissemination profiles play distinct roles in amplifying content. On Telegram, we observed that information flow is characterized by fragmented structures, with message dissemination occurring within smaller, ideologically aligned communities. This structure suggests that the platform is geared more toward strategic mobilization within closed groups rather than widespread, viral sharing. By incorporating topic modeling, we also showed that these dissemination profiles are linked to specific types of content, clarifying how certain narratives gain prominence depending on the underlying network structure.

This study has several limitations that should be acknowledged. First, our model is based on co-interaction heuristics, which might not identify more subtle or indirect coordination methods. Second, our unsupervised topic modeling involves a degree of subjective interpretation and depends on parameter tuning without a ground-truth reference. Finally, the dissemination patterns we observed are specific to this context and may not be generalizable to other platforms or political events.

These limitations highlight promising directions for future research. One potential avenue is to explore alternative metrics for temporal aggregation and different ways to define interaction speed. Another opportunity involves broadening the textual analysis to include sentiment, toxicity, or stance detection, which could provide richer interpretations of the dissemination profiles. Future studies could also investigate how certain dissemination classes amplify content from specific influential or ideological sources. Lastly, applying this framework to other platforms or to real-time monitoring systems could help track information dynamics as they unfold, aiding in the development of strategies to counter disinformation in digital environments.

Declarations

Acknowledgements

The authors would like to thank the colleagues and collaborators who contributed valuable discussions and insights throughout the development of this work. We are also grateful to *30° Simpósio Brasileiro de Sistemas Multimídia e Web* reviewers for their constructive feedback that helped to improve this investigation. In addition, we thank the reviewers and editors of the Journal on Interactive Systems for the opportunity to present and refine our work. We also acknowledge the support and collaboration of the *DDL Lab* at

UFOP and the *Locus Lab* at UFMG, whose environments fostered enriching academic exchange and inspiration for this research.

Funding

The Universidade Federal de Ouro Preto financially supported this study through the Scientific and Technological Initiation Scholarship Program, the *Conselho Nacional de Desenvolvimento Científico e Tecnológico* (CNPq) and the Universidade Federal de São João del-Rei. R.S.F acknowledges the financial support by the *Fundação de Amparo à Pesquisa do Estado de Minas Gerais* (FAPEMIG)-Brazil and *Conselho Nacional de Desenvolvimento Científico e Tecnológico* (CNPq)-Brazil (Grant No. BPQ-100-24/APQ-06611-24). J.M.A also acknowledges the financial support by the *Conselho Nacional de Desenvolvimento Científico e Tecnológico* (CNPq) and the *Fundação de Amparo à Pesquisa do Estado de Minas Gerais* (FAPEMIG), Brazil. Last, we thank the *Instituto Nacional de Ciência e Tecnologia em Inteligência Artiicial Responsável para Linguística Computacional, Tratamento e Disseminação de Informação* (INCT-TILD-IAR) for its support.

Authors' Contributions

J.M.A., A.P.C.S., R.S.F., and C.H.G.F led the conceptualization and methodology. The investigation and data analysis were performed by GSO, PJL, and OV, who were also responsible for the software implementation. GSO and JPL wrote the original draft, while JMA, APCS, RSF, and CGF contributed to the review and editing. J.M.A., A.P.C.S., R.S.F., and C.H.G.F. supervised and administrated the project.

Competing interests

The authors declare that they have no competing interests.

Availability of data and materials

All materials used in this study are openly available at: `https://github.com/gseovana/bidimensional-network-backbone-framework`.
This repository is provided to facilitate consultation, ensure reproducibility, and promote the principles of open science and knowledge sharing. All materials are distributed under the MIT license.

References

Abilov, A., Hua, Y., Matatov, H., Amir, O., and Naaman, M. (2021). Voterfraud2020: a multi-modal dataset of election fraud claims on twitter. In *International Conference on Web and Social Media*. DOI: http://dx.doi.org/10.1609/icwsm.v15i1.18113.

Albarrak, M. S., Elnahass, M., Papagiannidis, S., and Salama, A. (2020). The effect of twitter dissemination on cost of equity: A big data approach. *International Journal of Information Management*, 50:1–16. DOI: https://doi.org/10.1016/j.ijinfomgt.2019.04.014.

Antonakaki, D., Fragopoulou, P., and Ioannidis, S. (2021). A survey of twitter research: Data model, graph structure, sentiment analysis and attacks. *Expert systems with applications*, 164:114006. DOI: https://doi.org/10.1016/j.eswa.2020.114006.

Araujo, M. M., Ferreira, C. H., Reis, J. C., Silva, A. P., and Almeida, J. M. (2023). Identificação e caracterização de campanhas de propagandas eleitorais antecipadas brasileiras no twitter. In *Anais do XII Brazilian Workshop on Social Network Analysis and Mining*, pages 67–78. SBC. DOI: https://doi.org/10.5753/brasnam.2023.229879.

Armstrong, G., Martino, C., Rahman, G., Gonzalez, A., Vázquez-Baeza, Y., Mishne, G., and Knight, R. (2021). Uniform manifold approximation and projection (umap) reveals composite patterns and resolves visualization artifacts in microbiome data. *MSystems*, 6(5):10–1128. DOI: https://doi.org/10.1128/msystems.00691-21.

Badawy, A., Addawood, A., Lerman, K., and Ferrara, E. (2019). Characterizing the 2016 russian ira influence campaign. *Social Network Analysis and Mining*, 9:1–11. DOI: https://doi.org/10.1007/s13278-019-0578-6.

Barabási, A.-L. (2013). Network science. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 371(1987):20120375. DOI: https://doi.org/10.1098/rsta.2012.0375.

Barbosa, C., Félix, L., Vieira, V., and Xavier, C. (2019). Sara - a semi-automatic framework for social network analysis. In *Anais Estendidos do XXV Simpósio Brasileiro de Sistemas Multimídia e Web*. DOI: https://doi.org/10.5753/webmedia_estendido.2019.8137.

Bellutta, D. and Carley, K. M. (2023). Investigating coordinated account creation using burst detection and network analysis. *Journal of big Data*, 10(1):20. DOI: https://doi.org/10.1186/s40537-023-00695-7.

Bizel, G. and Singh, A. K. (2023). Political polarization, misinformation, and sentiments: A social media analysis about 'capitol hill 2021 attack'by tweets. *Journal of Social, Humanities and Administrative Sciences*, 9(61):2257–2266. DOI: http://dx.doi.org/10.29228/JOSHAS.64755.

Blondel, V. D., Guillaume, J.-L., Lambiotte, R., and Lefebvre, E. (2008). Fast unfolding of comm. in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2008(10):P10008. DOI: https://doi.org/10.1088/1742-5468/2008/10/P10008.

Castagna, S., Porrino, G., and Borgonovo, F. (2023). The italian pro-russia digital ecosystem on telegram. *Cybersecurity and Law*, 10(2):299–317.

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. DOI: https://doi.org/10.1016/j.dss.2022.113819.

Coscia, M. and Neffke, F. M. (2017). Network backboning with noisy data. In *2017 IEEE*

33rd international conference on data engineering (ICDE), pages 425–436. IEEE. DOI: https://doi.ieeecomputersociety.org/10.1109/ICDE.2017.10.

da Rosa, J. M., Linhares, R. S., Ferreira, C. H. G., Nobre, G. P., Murai, F., and Almeida, J. M. (2022). Uncovering discussion groups on claims of election fraud from twitter. In *Proc. of Social Informatics: 13th International Conference*. DOI: https://doi.org/10.1007/978-3-031-19097-1_20.

Danaditya, A., Ng, L. H. X., and Carley, K. M. (2022). From curious hashtags to polarized effect: profiling coordinated actions in indonesian twitter discourse. *Social Network Analysis and Mining*, 12(1):105. DOI: https://doi.org/10.1007/s13278-022-00936-2.

Dargahi Nobari, A., Reshadatmand, N., and Neshati, M. (2017). Analysis of telegram, an instant messaging service. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, CIKM '17. DOI: https://doi.org/10.1145/3132847.3133132.

de Souza, F. C. (2020). *BERTimbau: pretrained BERT models for Brazilian Portuguese BERTimbau: modelos BERT pré-treinados para Português Brasileiro*. PhD thesis, [sn].

Elmas, T., Overdorf, R., Özkalay, A. F., and Aberer, K. (2021). Ephemeral astroturfing attacks: The case of fake twitter trends. In *2021 IEEE European symposium on security and privacy (EuroS&P)*, pages 403–422. IEEE. DOI: https://doi.org/10.1109/EuroSP51992.2021.00035.

Ferreira, C., Murai, F., Silva, A., Almeida, J., Trevisan, M., Vassio, L., Mellia, M., and Drago, I. (2021). On the dynamics of political discussions on instagram: A network perspective. *Online Social Networks and Media*. DOI: https://doi.org/10.1016/j.osnem.2021.100155.

Ferreira, C. H., Murai, F., Matos, B., and Almeida, J. M. (2019). Modeling dynamic ideological behavior in political networks. *Web Science Journal*, 7. DOI: https://doi.org/10.34962/jws-80.

Ferreira, C. H. G., Murai, F., Couto da Silva, A. P., de Almeida, J. M., Trevisan, M., Vassio, L., Drago, I., and Mellia, M. (2020). Unveiling community dynamics on instagram political network. In *12th ACM Conference on Web Science*. DOI: https://doi.org/10.1145/3394231.3397913.

Fortunato, S. (2010). Community detection in graphs. *Physics reports*, 486(3-5):75–174. DOI: https://doi.org/10.1016/j.physrep.2009.11.002.

Fortunato, S. and Hric, D. (2016). Community detection in networks: A user guide. *Physics reports*, 659:1–44. DOI: https://doi.org/10.1016/j.physrep.2016.09.002.

G. Da Fonseca, L. G., Gomes Ferreira, C. H., and Soares Dos Reis, J. C. (2024). The role of news source certification in shaping tweet content: Textual and dissemination patterns in brazil's 2022 elections. In *Proceedings of the 20th Brazilian Symposium on Information Systems*, pages 1–10. DOI: https://doi.org/10.1145/3658271.3658303.

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: https://doi.org/10.1080/1369118X.2020.1739732.

Golovchenko, Y., Buntain, C., Eady, G., Brown, M. A., and Tucker, J. A. (2020). Cross-platform state propaganda: Russian trolls on twitter and youtube during the 2016 us presidential election. *The International Journal of Press/Politics*, 25(3):357–389. DOI: https://doi.org/10.1177/1940161220912682.

Gomes Ferreira, C. H., Murai, F., Silva, A. P. C., Trevisan, M., Vassio, L., Drago, I., Mellia, M., and Almeida, J. M. (2022). On network backbone extraction for modeling online collective behavior. *PLOS ONE*, 17(9):1–36. DOI: https://doi.org/10.1371/journal.pone.0274218.

Grady, D., Thiemann, C., and Brockmann, D. (2012). Robust classification of salient links in complex networks. *Nature communications*, 3(1):864. DOI: https://doi.org/10.1038/ncomms1847.

Graham, T., Hames, S., and Alpert, E. (2024). The coordination network toolkit: a framework for detecting and analysing coordinated behaviour on social media. *Journal of Computational Social Science*, pages 1–22. DOI: https://doi.org/10.1007/s42001-024-00260-z.

Grimminger, L. and Klinger, R. (2021). Hate towards the political opponent: A twitter corpus study of the 2020 us elections on the basis of offensive speech and stance detection. In *Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*.

Grootendorst, M. (2022). Bertopic: Neural topic modeling with a class-based tf-idf procedure. *arXiv preprint arXiv:2203.05794*. DOI: http://dx.doi.org/10.48550/arXiv.2203.05794.

Gunawan, D., Harahap, S. H., and Rahmat, R. F. (2019). Multi-document summarization by using textrank and maximal marginal relevance for text in bahasa indonesia. In *2019 International conference on ICT for smart society (ICISS)*, volume 7, pages 1–5. IEEE. DOI: https://doi.org/10.1109/ICISS48059.2019.8969785.

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. DOI: https://doi.org/10.1080/10584609.2019.1661888.

Khan, B. S. and Niazi, M. A. (2017). Network community detection: A review and visual survey. *arXiv preprint arXiv:1708.00977*. DOI: https://doi.org/10.48550/arXiv.1708.00977.

Lanius, C., Weber, R., and MacKenzie Jr, W. I. (2021). Use of bot and content flags to limit the spread of misinformation among social networks: a behavior and attitude survey. *Social network analysis and mining*, 11(1):32. DOI: https://doi.org/10.1007/s13278-021-00739-x.

Lima, F. M. d. C., de Miranda, L. C., Vasiljevic, G. A. M., and Baranauskas, M. C. C. (2024). An analysis of the authorship and co-authorship networks of the brazilian human-computer interaction conference. *Journal on Interactive Systems*, 15(1):265–293. DOI: https://doi.org/10.5753/jis.2024.3340.

Linhares, R. S., Rosa, J. M., Ferreira, C. H. G., Murai, F., Nobre, G., and Almeida, J. (2022). Uncovering coordinated communities on twitter during the 2020 u.s. election. In *Proc. of ASONAM*. DOI:

https://doi.org/10.1109/ASONAM55673.2022.10068628.

Malagoli, L., Piorino, G., Ferreira, C. H., and da Silva, A. P. C. (2024). Twitter and the 2022 brazilian elections portrait: A network and content-driven analysis. In *Brazilian Symposium on Multimedia and the Web (WebMedia)*, pages 283–291. SBC. DOI: https://doi.org/10.5753/webmedia.2024.241926.

Mao, Y., Qu, Y., Xie, Y., Ren, X., and Han, J. (2020). Multi-document summarization with maximal marginal relevance-guided reinforcement learning. *arXiv preprint arXiv:2010.00117*. DOI: https://doi.org/10.18653/v1/2020.emnlp-main.136.

Marcaccioli, R. and Livan, G. (2019). A pólya urn approach to information filtering in complex networks. *Nature communications*, 10(1):745. DOI: https://doi.org/10.1038/s41467-019-08667-3.

Markopoulou, A., Iannaccone, G., Bhattacharyya, S., Chuah, C.-N., Ganjali, Y., and Diot, C. (2008). Characterization of failures in an operational ip backbone network. *IEEE/ACM transactions on networking*, 16(4):749–762. DOI: https://doi.org/10.1109/TNET.2007.902727.

Neal, Z. P. (2022). backbone: An r package to extract network backbones. *PloS one*, 17(5):e0269137. DOI: https://doi.org/10.1371/journal.pone.0269137.

Neal, Z. P., Domagalski, R., and Sagan, B. (2021). Comparing alternatives to the fixed degree sequence model for extracting the backbone of bipartite projections. *Scientific reports*, 11(1):23929. DOI: https://doi.org/10.1038/s41598-021-03238-3.

Neal, Z. P. and Neal, J. W. (2024). Illustrating the importance of edge constraints in backbones of bipartite projections. *Plos one*, 19(5):e0302973. DOI: https://doi.org/10.1371/journal.pone.0302973.

Newman, M. (2018). Network structure from rich but noisy data. *Nature Physics*, 14:542–546. DOI: https://doi.org/10.1038/s41567-018-0076-1.

Newman, M. E. (2006). Finding community structure in networks using the eigenvectors of matrices. *Physical review E*, 74(3):036104. DOI: https://doi.org/10.1103/PhysRevE.74.036104.

Newman, M. E. and Girvan, M. (2004). Finding and evaluating community structure in networks. *Physical review E*, 69:026113. DOI: https://doi.org/10.1103/PhysRevE.69.026113.

Ng, L. H. X., Kloo, I., Clark, S., and Carley, K. M. (2024). An exploratory analysis of covid bot vs human disinformation dissemination stemming from the disinformation dozen on telegram. *Journal of Computational Social Science*, pages 1–26. DOI: https://doi.org/10.1007/s42001-024-00253-y.

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. DOI: https://doi.org/10.1609/icwsm.v15i1.18074.

Nobre, G., Ferreira, C., and Almeida, J. (2020). Beyond groups: Uncovering dynamic communities on the whatsapp network of information dissemination. In *SocInfo' 2020*. DOI: https://doi.org/10.1007/978-3-030-60975-7_19.

Nobre, G. P., Ferreira, C. H., and Almeida, J. M. (2022). A hierarchical network-oriented analysis of user participation in misinformation spread on whatsapp. *Information Processing and Management*, 59(1). DOI: https://doi.org/10.1016/j.ipm.2021.102757.

Oliveira, G. S., Venâncio, O., Vieira, V., Almeida, J., Silva, A. P., Ferreira, R., and Ferreira, C. H. (2024). Um framework para análise bidimensional de disseminação de informações em plataformas de mídias sociais. In *Proceedings of the 30th Brazilian Symposium on Multimedia and the Web (WebMedia 2024)*. Sociedade Brasileira de Computação-SBC. DOI: https://doi.org/10.5753/webmedia.2024.241957.

Pacheco, D., Flammini, A., and Menczer, F. (2020). Unveiling coordinated groups behind white helmets disinformation. In *The Web Conference*. DOI: https://doi.org/10.1145/3366424.3385775.

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 *International Conference on Web and Social Media*. DOI: https://doi.org/10.1609/icwsm.v15i1.18075.

Pohl, J. S., Markmann, S., Assenmacher, D., and Grimme, C. (2023). Invasion@ ukraine: providing and describing a twitter streaming dataset that captures the outbreak of war between russia and ukraine in 2022. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 17, pages 1093–1101. DOI: https://doi.org/10.1609/icwsm.v17i1.22217.

Reimers, N. (2019). Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084*. DOI: https://doi.org/10.18653/v1/D19-1410.

Rossetti, G., Pappalardo, L., Pedreschi, D., and Giannotti, F. (2017). Tiles: an online algorithm for community discovery in dynamic social networks. *Machine Learning*, 106:1213–1241. DOI: https://doi.org/10.1007/s10994-016-5582-8.

Rossini, P., Mont'Alverne, C., and Kalogeropoulos, A. (2023). Explaining beliefs in electoral misinformation in the 2022 brazilian election: The role of ideology, political trust, social media, and messaging apps. *Harvard Kennedy School Misinformation Review*, 4(3). DOI: https://doi.org/10.37016/mr-2020-115.

Saxena, N., Sinha, A., Bansal, T., and Wadhwa, A. (2023). A statistical approach for reducing misinformation propagation on twitter social media. *Information Processing & Management*, 60(4):103360. DOI: https://doi.org/10.1016/j.ipm.2023.103360.

Serrano, M. Á., Boguná, M., and Vespignani, A. (2009). Extracting the multiscale backbone of complex weighted networks. *Proceedings of the National Academy of Sciences*, 106(16):6483–6488. DOI: https://doi.org/10.1073/pnas.0808904106.

Sharma, K., Zhang, Y., Ferrara, E., and Liu, Y. (2021). Identifying coordinated accounts on social media through hidden influence and group behaviours. In *Proceed-*

ings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining, pages 1441–1451. DOI: https://doi.org/10.1145/3447548.3467391.

Silva, F. and Freitas, L. (2022). Brazilian portuguese hate speech classification using bertimbau. In *The International FLAIRS Conference Proceedings*, volume 35. DOI: https://doi.org/10.32473/flairs.v35i.130594.

Souravlas, S., Sifaleras, A., Tsintogianni, M., and Katsavounis, S. (2021). A classification of community detection methods in social networks: a survey. *International Journal of General Systems*, 50(1):63–91. DOI: https://doi.org/10.1080/03081079.2020.1863394.

Souza, F., Nogueira, R., and Lotufo, R. (2020). Bertimbau: Pretrained bert models for brazilian portuguese. In Cerri, R. and Prati, R. C., editors, *Intelligent Systems*, pages 403–417, Cham. Springer International Publishing. DOI: https://doi.org/10.1007/978-3-030-61377-8_28.

Stieglitz, S. and Dang-Xuan, L. (2013). Social media and political communication: a social media analytics framework. *Social network analysis and mining*, 3:1277–1291. DOI: https://doi.org/10.1007/s13278-012-0079-3.

Traag, V. A., Waltman, L., and Van Eck, N. J. (2019). From louvain to leiden: guaranteeing well-connected communities. *Scientific reports*, 9(1):1–12. DOI: https://doi.org/10.1038/s41598-019-41695-z.

Ulizko, M., Artamonov, A., Tukumbetova, R., Antonov, E., and Vasilev, M. (2022). Critical paths of information dissemination in networks. *Scientific Visualization*, 14(2):98–107. DOI: https://doi.org/10.26583/sv.14.2.09.

Urman, A. and Katz, S. (2020). What they do in the shadows: examining the far-right networks on telegram. *Information, Communication & Society*, pages 1–20. DOI: https://doi.org/10.1080/1369118X.2020.1803946.

Vargas, L., Emami, P., and Traynor, P. (2020). On the detection of disinformation campaign activity with network analysis. In *ACM Cloud Computing Security Workshop*, pages 133–146. DOI: https://doi.org/10.1145/3411495.3421363.

Varshney, D. and Vishwakarma, D. K. (2021). A review on rumour prediction and veracity assessment in online social network. *Expert Systems with Applications*, 168:114208. DOI: https://doi.org/10.1016/j.eswa.2020.114208.

Venâncio, O. R., Ferreira, C. H., Almeida, J. M., and da Silva, A. P. C. (2024a). Unraveling user coordination on telegram: A comprehensive analysis of political mobilization during the 2022 brazilian presidential election. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 18, pages 1545–1556. DOI: https://doi.org/10.1609/icwsm.v18i1.31408.

Venâncio, O. R., Gonçalves, G. H., Ferreira, C. H., and da Silva, A. P. C. (2024b). Evidências de disseminação de conteúdo no telegram durante o ataque aos órgãos públicos brasileiros em 2023. In *Brazilian Symposium on Multimedia and the Web (WebMedia)*, pages 385–389. SBC. DOI: https://doi.org/10.5753/webmedia.2024.241972.

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. In *Proceedings of the international AAAI conference on web and social media*, volume 18, pages 1557–1570. DOI: http://dx.doi.org/10.1609/icwsm.v18i1.31409.

Weber, D. and Neumann, F. (2020). Who's in the gang? revealing coordinating communities in social media. In *2020 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pages 89–93. IEEE. DOI: https://doi.org/10.1109/asonam49781.2020.9381418.

Weber, D. and Neumann, F. (2021). Amplifying influence through coordinated behaviour in social networks. *Social Network Analysis and Mining*, 11(1):111. DOI: https://doi.org/10.1007/s13278-021-00815-2.

Zayats, V. and Ostendorf, M. (2018). Conversation modeling on reddit using a graph-structured lstm. *Transactions of the Association for Computational Linguistics*, 6:121–132. DOI: https://doi.org/10.1162/tacl_a_00009.

Zhuravskaya, E., Petrova, M., and Enikolopov, R. (2020). Political effects of the internet and social media. *Annual review of economics*, 12(1):415–438. DOI: https://doi.org/10.1146/annurev-economics-081919-050239.