International Journal of Public Opinion Research, **36**, edae032
https://doi.org/10.1093/ijpor/edae032
Advance access publication 08 July, 2024
Article

OXFORD

Analyzing Misinformation Claims During the 2022 Brazilian General Election on WhatsApp, Twitter, and Kwai

Scott A. Hale[1,2], **Adriano Belisario**[1], **Ahmed Nasser Mostafa**[2], **Chico Camargo**[1,3,4]

[1]Oxford Internet Institute, University of Oxford, Oxford, UK
[2]Meedan, Inc., San Francisco, USA
[3]Department of Computer Science, University of Exeter, UK
[4]Division of English Language & Literature, Ewha Womans University, Seoul, South Korea

All correspondence concerning this article should be addressed to: Scott A. Hale, Oxford Internet Institute, University of Oxford, 1 St Giles, Oxford, UK.
Email: scott@meedan.com

Abstract

This study analyzes misinformation claims sent to fact-checking organizations on WhatsApp during the 2022 Brazilian general election and compares them with content from Twitter and Kwai (a popular video-sharing application similar to TikTok). Given the democratic importance of accurate information during elections, multiple fact-checking organizations collaborated to collect and respond to misinformation via WhatsApp tiplines and power a fact-checking feature within a chatbot operated by Brazil's election authority, the Tribunal Superior Eleitoral (TSE). We partnered with TSE and three fact-checking organizations and collected social media data to study how misinformation claims propagate across platforms. We observed little overlap between the users of different fact-checking tiplines and a high correlation between the number of users and the amount of unique content, suggesting that WhatsApp tiplines are far from reaching a saturation point. Similarly, we also found little overlap in content across platforms, indicating the need for further research with cross-platform approaches to identify misinformation dynamics.

Introduction

User-generated content platforms are shaped by their users and designs. Adoption of platforms is uneven across demographics (Jungherr, 2018; Ruths & Pfeffer, 2014), and this can easily lead to unique content on each platform. The unique designs of platforms also signal different affordances (Gibson, 1979), which could lead to unique content. For example, it is clear that Instagram favors visual content, Facebook favors longer-form text, and Twitter favors shorter-form content. While many studies of online misinformation focus on one platform—often Twitter (now X), which has generally been over-researched to the detriment of other platforms (Cihon & Yasseri, 2016)—the spread of misinformation is not constrained by platform boundaries. Instead, it flows and transforms across different digital environments, influenced by platform-specific affordances, user behaviors, and social contexts (Litt, 2012; Norman, 2013).

Many people use multiple platforms in their everyday lives (Blank, Dutton, & Lefkowitz, 2020), and these users are embedded in various environments and social contexts with a range of interactions with other individuals both online and offline (Lamb, King, & Kling, 2003). Using multiple platforms provides a possible mechanism for the same content to spread across platforms. In some settings, for example, it is common to find screenshots of Twitter posts on Instagram (Asian American Disinformation Table, 2022). These different dynamics underscore the importance of investigating misinformation not just within single platforms but also in its movement and mutation across different platforms.

As the fourth-largest democracy in the world and a country with significant Internet penetration (Cetic.br, 2021), the Brazilian general elections present a good opportunity to investigate political misinformation at scale. We partnered with Brazil's election authority, the Tribunal Superior Eleitoral (TSE), and fact-checking organizations to obtain misinformation claims related to the 2022 general elections and collected social media data from Twitter (now X) and Kwai (a popular video-sharing app similar to TikTok). In this article, we focus specifically on *potential misinformation claims*, which we operationalize as content citizens submit to chatbot tiplines run by fact-checking organizations or the TSE. Instead of investigating the veracity of claims ourselves, our primary goal is to perform a descriptive analysis of the claims and measure how unique the claims circulating on each platform are.

To the best of our knowledge, this article is the first large-scale analysis of misinformation using data from Kwai, Twitter, and WhatsApp. Overall, we find that the unique characteristics of each platform influence the specific content circulating. These differences complicate large-scale quantitative comparisons, and this article highlights areas where further algorithmic development is needed. While there is overlap across the platforms and the heavy-tailed patterns often observed in collective human behavior (Margetts, John, Hale, & Yasseri, 2015) can be observed here, we also find clear differences between platforms.

Related Work

Cross-platform Misinformation

Scholarship about "misinformation" online appears only a few years after the first graphical web browser: Hernon (1995) defines disinformation as a "deliberate attempt to deceive or mislead" and misinformation as "an honest mistake" (p. 134). As Internet use has grown, so too has scholarship on the topic (Ha, Andreu Perez, & Ray, 2021), and various alternative definitions have been proposed. Similar to other scholars, we use "misinformation" as an umbrella term for any misleading content regardless of the author's intentions (see, e.g., Pantazi, Hale, & Klein, 2021).[1]

Misinformation is not a problem exclusive to the Internet (Altay, Berriche, & Acerbi, 2023). There are a variety of psychological processes at play in the spread of misinformation including excess gullibility and excess vigilance (Pantazi, Hale, & Klein, 2021). At the same time, motivated reasoning or directional motives—that is, the desire to arrive at a specific conclusion—also play a role, specifically in political misinformation (Jerit & Zhao, 2020).

Similarly, as people often access multiple platforms and are embedded across a range of environments and social contexts (Lamb, King, & Kling, 2003), scholars also emphasize the role of group identities and moral emotions in shaping political beliefs and behaviors online (Pereira, Harris, & Van Bavel, 2023; Rathje, Van Bavel, & Van Der Linden, 2021; Van Bavel, Rathje, Vlasceanu, & Pretus, 2024). This perspective suggests that people are not only influenced by platform affordances but also by their social identities and emotional engagements with content. For instance, partisan alignment can intensify misinformation sharing, driven by loyalty to group narratives rather than platform-specific features (Osmundsen, Bor, Vahlstrup, Bechmann, & Petersen, 2021).

Moreover, the design of social media platforms can also influence the spread of misinformation as design can influence how much people think about the accuracy of the content they are sharing (Pennycook et al., 2021). More broadly, the design of user interfaces influences what actions users perceive are available for them and the intended uses of a platform (Gibson, 1979; Norman, 2013).

All these factors form a complex web of interactions. We argue that while the design of social media platforms guides user actions and perceived platform purposes (Gibson, 1979; Norman, 2013; Pennycook et al., 2021), it is the interplay of these designs with users' social identities and pre-existing motivations (Osmundsen, Bor, Vahlstrup, Bechmann, & Petersen, 2021; Pereira, Harris, & Van Bavel, 2023) that shapes the information sharing landscape. A comprehensive analysis of this phenomenon must therefore consider the complex interplay of affordance theory, social identity, emotional engagement, and group dynamics.

The concept of imagined audiences (Litt, 2012) and the interplay of social identity and platform design (Pennycook et al., 2021; Pereira, Harris, & Van Bavel, 2023) suggest that while the core of a misinformation claim might remain constant, its presentation, emphasis, and framing could vary significantly across platforms. These variances are pivotal in understanding how misinformation resonates and spreads within different communities.

Recent work on cross-platform content sharing describes media ecosystems—or "echo-systems" (Starbird et al., 2018)—as rife with a large degree of content overlap between platforms, but also with content being adapted, shared, and remixed across platforms, not only organically, but often as part of disinformation campaigns—something that is often hard to quantify due to the many forms in which a piece of content might be presented (Starbird et al., 2018; Thorson et al., 2013; Wilson & Starbird, 2020). This cross-talk between different digital spaces demands a more nuanced approach to studying misinformation, one that encompasses the entire digital ecosystem rather than isolated platforms.

Social Media Platforms and the 2022 Brazilian Election

Our study is set during the 2022 Brazilian election. The first round of voting was held on 2 October and included elections for the houses of the Brazilian Congress, state governors, and the president. After the first round, no presidential candidate captured over 50% of the ballots cast, and a run-off election was conducted on 30 October between the top two candidates: Luiz Inácio Lula da Silva (former president of Brazil, from 2003 to 2010) and Jair Bolsonaro (elected in 2018, seeking re-election). Ultimately, Lula da Silva was elected for his third mandate, defeating the incumbent Bolsonaro. Similar to many recent democratic elections, the run-up to the Brazilian general elections was rife with claims of illegitimacy of the electoral process, as well as misleading information regarding multiple aspects of Brazilian politics (Rossini, Mont'Alverne, & Kalogeropoulos, 2023; Tarouco, 2023).

Social media platforms are a critical part of contemporary Brazilian politics, serving as both a catalyst and a conduit for political discourse, activism, and misinformation (Dwoskin, 2023; Reis, Melo, Garimella, & Benevenuto, 2020; Reis, Garimella, Almeida, Eckles, & Benevenuto, 2020; Resende et al., 2019). Of the Brazilian population aged 10 or older, 81% were Internet users in 2021 (Cetic.br, 2021) and WhatsApp is by far the most used platform in Brazil (Paiva, 2020): It is installed in over 99% of smartphones in Brazil (Paiva, 2020) and 57% of Brazilians are using it for news consumption or as a source of political information (Anita Baptista, Rossini, Veiga de Oliveira, & Stromer-Galley , 2019; Newman, Fletcher, Robertson, Eddy, & Nielsen, 2022). In addition to WhatsApp, surveys show that Twitter and Kwai have considerable userbases in the country (Gava, 2022; Kemp, 2022; Paiva, 2022).

Concerned with the spread of online misinformation during the electoral period, Brazil's TSE signed agreements with WhatsApp, Twitter, Kwai, and other major platforms to fight misleading content about the electoral processes of the 2022 general elections. We briefly contextualize each social platform analyzed.

[1] We include studies of "fake news" but refrain from using the term here given its imprecision and politicization (Brummette, DiStaso, Vafeiadis, & Messner, 2018).

WhatsApp is Brazil's most popular application and the second most popular app for news consumption (Newman, Fletcher, Robertson, Eddy, & Nielsen, 2022). Previous studies have documented the role of WhatsApp groups and bulk messages for misinformation networks in the 2018 general elections (Machado, Kira, Narayanan, Kollanyi, & Howard, 2019; Resende et al., 2019). Brazilian media outlets and government authorities have responded to misinformation spread on WhatsApp by creating their own chatbot tiplines to receive and fact-check claims submitted by citizens. These tiplines are accounts on WhatsApp to which users can send possible misinformation content and questions and, in exchange, receive fact-checks and trusted information (Johansson et al., 2022). In 2022, the TSE partnered with WhatsApp for the 2022 elections. The partnership included WhatsApp's support in developing a chatbot tipline for this purpose (TSE, 2022). The WhatsApp data we analyze come from anonymous questions submitted via that fact-checking feature and similar tiplines implemented by three Brazilian fact-checking initiatives.

Kwai is owned by Kuaishou Technology, a Chinese company which operates a short video platform known as Snack Video in South Asia and Kwai in Latin America. Kuaishou's platforms now have more than one billion users globally. The company opened its Brazilian branch in 2019 and claimed to have 45 million monthly active users in Brazil during the 2022 general election, most located in less economically privileged regions in Brazil (Deck & Marasciulo, 2022). We are unaware of any prior peer-reviewed scholarship on Kwai.

Twitter (now X) had about 19 million estimated users in Brazil during the 2022 elections (Kemp, 2022). Although not as popular as WhatsApp and Kwai, Twitter has a number of important features. The platform is popular among key actors with influence on public opinion, including politicians, journalists, researchers, activists, and company executives. Twitter was intensively used for political debate in the 2018 general elections, and the role of bots or automated accounts in misinformation networks has been documented by researchers (Recuero, 2020). In its agreement with the TSE for the 2022 elections, Twitter agreed to activate search prompts and warnings with official information on top of contents related to the elections, create special content, and provide follow-up on misinformation denunciations made by the TSE.

Research Questions

We collected data from Twitter and Kwai and partnered with Brazil's election authority, the TSE, and three fact-checking organizations to collect data from WhatsApp chatbot tiplines. These tiplines were designed to receive misinformation claims from citizens and facilitate fact-checking. We perform a descriptive analysis of the datasets—examining the formats and content of popular misinformation claims—and compare the overlap of these claims among WhatsApp tiplines and these two major social platforms.

We investigate two research questions:

RQ1 To what extent does information overlap between fact-checker WhatsApp tiplines and the WhatsApp bot provided by Brazil's election authority during the 2022 Brazilian elections?

RQ2 What overlap is there between misinformation sent to fact-checker WhatsApp tiplines with content found on Twitter and Kwai during the 2022 Brazilian elections?

Each organization we studied runs its own tipline, and each has its own audience and unique organizational environment (Lelo, 2022). Examining the overlap between the content sent to different fact-checkers or to the TSE bot (RQ1) will help us to understand how homogenous content is across different tiplines. If the content is quite diverse, it will suggest that each tipline captures only a small proportion of possible misinformation on WhatsApp. We also extend this exploration to a cross-platform context (RQ2), examining the information overlap across WhatsApp, Twitter, and Kwai during the 2022 Brazilian elections. Given the national focus on the election in Brazil and the fact that many people use multiple platforms, it is possible for similar misinformation content to exist on all three of our platforms. Analysis of public group and tipline data in the 2019 Indian elections found tiplines capture a significant proportion of popular content and identify that content quickly—often before it spreads in large groups (Kazemi, Garimella, Shahi, Gaffney, & Hale, 2022).

On the one hand, affordance theory suggests each platform will have unique characteristics and contain unique misinformation based on the respective affordances. On the other hand, the fact that many people use multiple platforms in their everyday lives points provides a mechanism for content to spread from one platform to another, suggesting a higher degree of overlap. The answers to these questions have important practical implications for misinformation response: If misinformation is relatively homogeneous across different platforms, then fact-checking organizations (and researchers) can monitor whatever platforms are most accessible. In contrast, if misinformation is relatively unique per platform, there is a greater need for gathering content in platform-specific ways.

Data and Methods

Table 1 shows a summary of the data sources. Twitter and WhatsApp posts include text, images, and videos, whereas Kwai posts comprise only videos and text for the video descriptions.

Table 1. Data Collection Methods and Statistics for Social Media Platforms

Method	Date range (inclusive), in 2022	No. of posts	Users
Twitter	September 20, 2022–November 10, 2022†	53,831,265	4,217,513
Kwai	October 8, 2022–November 26, 2022	23,737	13,068
WhatsApp (tiplines)	September 1, 2022–November 15, 2022	49,422	14,959
WhatsApp (TSE bot)	September 1, 2022–November 15, 2022	223,621	Unknown‡

†Twitter data collected between September 20, 2022–October 6, 2022 and October 21, 2022–November 10, 2022.
‡The TSE bot was anonymous and hence the number of unique users is unknown.

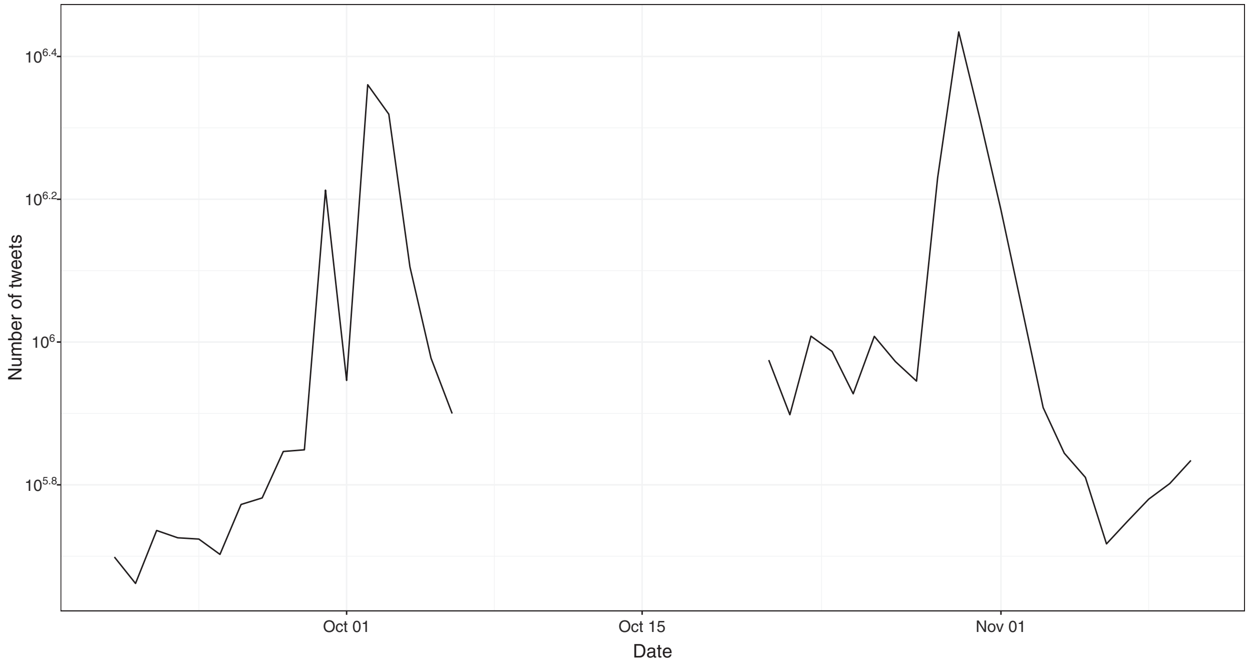

Figure 1. Volume of filtered, election-related tweets during the study period. Only original tweets in Portuguese are included (i.e., retweets are excluded).

We analyze anonymized WhatsApp data from three fact-checking organizations and the chatbot operated by the TSE. Our data include 49,422 submissions to fact-checker tiplines from 14,959 unique users as well as 223,621 submissions from an unknown number of TSE bot users between September 1 and November 15, 2022. All WhatsApp data are anonymous: Phone numbers were replaced with random IDs before we received the data and no other metadata beyond the timestamp were included in the data made available to us for analysis. Given the end-to-end encryption of WhatsApp, it is impossible to know exactly how representative any collection of messages is.

Although Kwai does not have an official API, we found a third-party service with APIs for the platform.[2] We used this service and validated it by comparing its data in real-time with the search results provided by queries run by a Kwai user in Brazil on October 23, 2022, using the same keywords. We found that the top videos suggested and their metrics were identical. We ran queries to this API with our location set to Brazil between October 8 and November 26, 2022. We added search keywords incrementally and found the results were sensitive to accents and capitalization. To account for this, we created variations for some of the search terms. We collected 589,878 search results and group them by user and video identification variables. This means that videos published by distinct usernames count as different posts. The search results corresponded to 35,701 unique videos. After preprocessing to remove empty and duplicate video descriptions, we identify 15,017 video descriptions that were used to cluster textual data in Kwai.

We capture tweets from Twitter using elevated access to its streaming API, which allowed us to avoid any rate limits. We capture tweets using a list of terms related to Brazilian politics that were initially sourced from news, Wikipedia, and social media content about the elections at the start of September 2022. After one day of content was captured, we analyzed the frequencies of unigrams, bigrams, and trigrams to identify additional terms to track. In total, 128 terms were tracked from September 20 to October 6, 2022, and from October 21 to November 10, 2022. The period in between was not tracked due to a data outage. The terms matched between 460 thousand and 2.7 million tweets per day. We filter all tweets in a second pass and discard all retweets. We detect the language of each tweet by first removing URLs and mentions and then applying the compact language detector (CLDv3).[3] Tweets not in Portuguese are discarded from our dataset. This results in a total of 53.8 million tweets from 4.2 million unique users. We find an average of approximately 800,000 original tweets per day. The volume of tweets over time shows two clear peaks corresponding with the elections (Figure 1).

Vectorization and Clustering

We represent text contents from WhatsApp, Kwai, and Twitter as dense vectors using a sentence-transformers MPNet language model trained to produce semantic sentence embeddings.[4] This model produces similar embeddings (i.e., vectors) for content with similar meanings even if the content uses different words (Reimers & Gurevych, 2019).

We use these embeddings when comparing within each platform and across the platforms. We compare all items to each other using cosine similarity and group items that are very similar. Our goal is to only group items that are really making the same claim, and we expect many items will be in a cluster

[2] https://rapidapi.com

[3] https://github.com/google/cld3

[4] https://huggingface.co/sentence-transformers/paraphrase-multilingual-mpnet-base-v2

by themselves. Therefore, we use single-link hierarchical clustering, which allows us to set a similarity threshold for how similar two items need to be to group together. This contrasts with *k*-means clustering, where the number of clusters (*k*) is specified and items often end up in larger clusters. Based on a manual examination of our dataset and previous work on this field (e.g., Kazemi, Garimella, Gaffney, & Hale, 2021), we use a threshold of 0.875. This approach to text clustering has been used and validated in previous misinformation research on WhatsApp (e.g., Kazemi, Garimella, Gaffney, & Hale, 2021).

In our initial testing, we found that capitalization, punctuation, and missing diacritics negatively affected our ability to match content across platforms. In particular, fact-checks often contained correct capitalization and diacritics, while social media content did not. To overcome this and make content more comparable, we took the text of each post, removed URLs, lower-cased all words, removed all punctuation, and replaced accents and other diacritics with their closest ASCII equivalent (e.g., we replaced “á” with “a”).

We clustered similar images using PDQ image hashing and similar videos using Temporal Match Kernel (TMK) embeddings. Both algorithms were developed by Meta Research to identify similar multimedia content and are described in detail in a whitepaper available online.[5] PDQ is a perceptual hashing algorithm that can identify similar pictures even if they have different file format or minor alternations. We calculate a normalized Hamming distance and use one minus this distance as a similarity measure. In line with previous research and the Meta whitepaper, we use a threshold of 0.7 for clustering images. PDQ is also employed by Reis, Melo, Garimella, and Benevenuto (2020) in their study of image sharing on WhatsApp.

We used TMK embeddings to compare the videos received on WhatsApp tiplines (including those forwarded by the TSE tipline) to Kwai videos. TMK embeddings only consider the visual portion of the video and not the audio. In addition to the TMK embeddings, we use the related `tmk-clusterize` executable provided by Meta to identify clusters of videos. We use similarity thresholds of 0.7 for both the level-1 and level-2 thresholds in accordance with the whitepaper produced by Meta.[6] TMK is a C++ library, but we develop and release open-source Python bindings to make it easier to calculate TMK embeddings in Python.[7]

Analysis

WhatsApp Tiplines

To measure the information overlap between the three fact-checker WhatsApp tiplines (RQ1), we adopt the following metrics and tests. First, we count the frequency of messages according to their type (video, images, link, text, or audio), pseudonymous author IDs, and dates. We also measure the overlap in users between tiplines.

One challenge in crowdsourcing the identification of misinformation via tiplines is building an active user base that will forward dubious content to fact-checking organizations. The relationship between the size of a tipline audience and the amount of novel content is unclear: Having a larger audience may simply result in multiple submissions of the same claim rather than the identification of new misinformation claims. We investigate this relationship in three ways: (1) we measure the overlap in users between tiplines for different types of messages; (2) we analyze the cumulative distributive function (CDF) of the number of messages per user and the number of items per cluster (cluster sizes); and (3) we randomly re-order users and consider the amount of novel content we could find with different subsets of random users.

We also compare data from the fact-checkers’ tiplines to the messages received by the TSE. The comparison assesses the total, length, and type of messages in both data sources. We examine the overlap and correlation of messages in the TSE Tipline section.

Kwai Videos and Descriptions

To understand to what extent information sent to fact-checker WhatsApp tiplines exists on Kwai (RQ2), we first perform a descriptive analysis of the data collected from Kwai, describing overall characteristics and the most important semantic clusters for the text descriptions of videos that circulated in this platform during the election. Videos without descriptions were omitted for this part of the analysis. We also cluster the videos and examine their overlap with the videos sent to fact-checkers via WhatsApp.

Twitter Text

To understand to what extent information sent to fact-checker WhatsApp tiplines exists on Twitter (RQ2), we select a representative message from the largest four clusters of WhatsApp messages sent to fact-checker tiplines. Then, we embed these top WhatsApp messages and all tweets, and select any pairs with cosine similarity above 0.7.

Manually inspecting the Twitter messages revealed that the semantic matches were often noisy. The messages sent to the WhatsApp tiplines were often very long—in fact, they often exceeded the 280-character limit that Twitter imposes—and this may, in part, explain the poor quality of the matches between the two platforms. To overcome the noise in the results, we further filtered the close semantic matches using keywords and manual inspection.

Results

Comparison of WhatsApp Tiplines

Fact-checker Tiplines

We start by examining the 49,422 submissions made to the three Brazilian fact-checkers’ WhatsApp tiplines. Most items submitted to the tiplines are videos (16,604) and images (11,463), followed by hyperlinks (11,363), text messages (8,613), and audio messages (1,379). We exclude 1,914 short text messages of less than 5 characters for the remainder of the analysis.

The submissions were made by 14,959 unique users. The number of unique items submitted per tipline user appears to be heavy-tailed (Figure 2, left). While 54% of users each submitted only one piece of content, the most active tipline user submitted 299 content items. Overall, only 5% of users are responsible for 41% of the unique items submitted to the tiplines.

[5] https://github.com/facebook/ThreatExchange/blob/main/hashing/hashing.pdf

[6] We confirm that a more stringent threshold of 0.9 does not significantly change the results.

[7] https://github.com/meedan/tmkpy

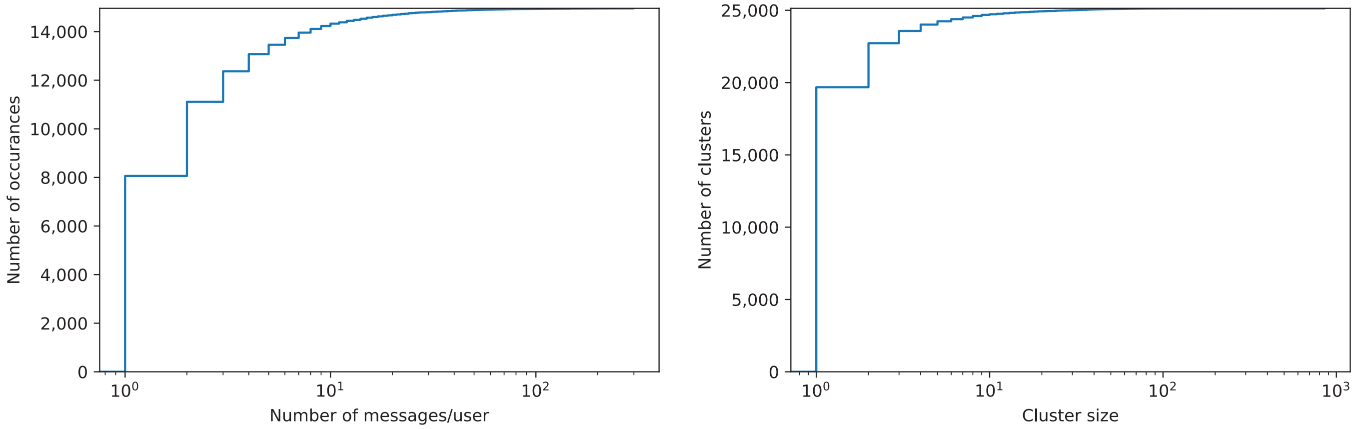

Figure 2. Left: The cumulative distribution function (CDF) comparing how many users (y-axis) submitted how many messages (x-axis) to the WhatsApp tiplines. While 54% of users sent only one message, the most prolific user sent 299. Right: The CDF comparing the number of clusters (y-axis) to their sizes (x-axis) on the WhatsApp tiplines. While 78% of clusters have only one item, the largest cluster has 858 instances of people thanking fact-checkers and the second largest has 210 instances of a video. Note that x-axes for both plots use log-10 scales.

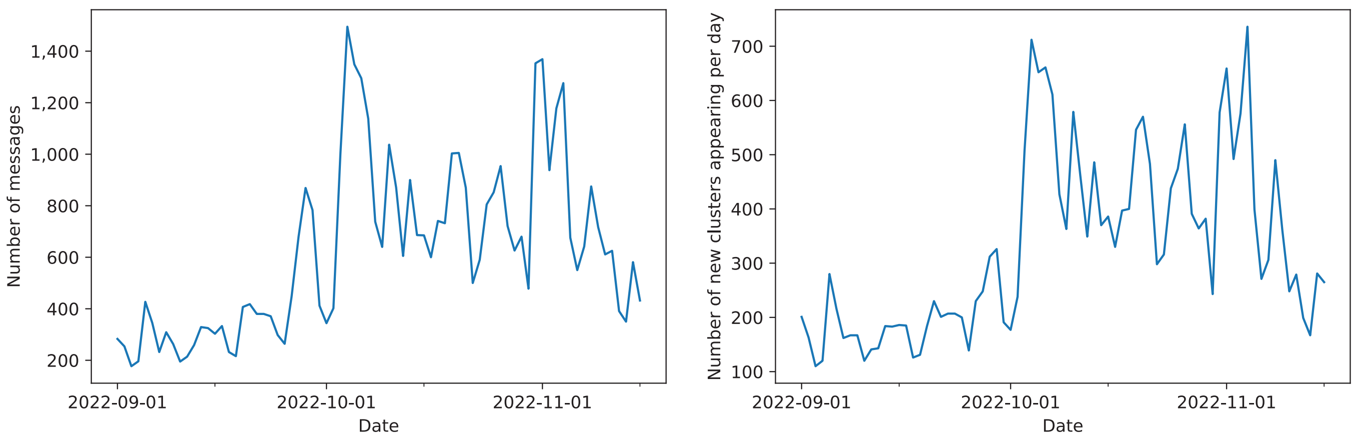

Figure 3. Left: Number of tipline submissions per day. Right: Number of new clusters appearing per day in the fact-checkers' misinformation tiplines.

We find the number of submissions per day to the tiplines varied considerably (Figure 3, left). There were pronounced peaks on the days of the two elections as well as another peak just after the run-off election. Manual inspection of content during this post-election peak revealed it was largely people inquiring about the results. The amount of novel content—represented by the number of new clusters (Figure 3, right)—follows a similar pattern to the number of daily messages, suggesting new messages often brought new content.

In general, we found little overlap between the users of different tiplines. Of the 6,383 users who submitted two or more messages, 93% submitted all their messages to the same tipline, 7% submitted at least one message to two tiplines, and 0.6% submitted messages to all three tiplines.

We find that novel textual content grows approximately linearly with the number of unique users (Figure 4). The empirically observed distribution is shown in blue, and 100 random re-orderings are shown in dashed pink lines. This relationship holds when considering all items (left) or only items that were fact-checked (right). The linear relationship suggests that more tipline users would likely lead to new content that would not otherwise be observed.

We manually examined the largest clusters of messages and found they were often video or long, forwarded text messages. We discarded a few large clusters that were greetings (e.g., "Bom dia"), expressions of gratitude to the fact-checkers, or spam.

The largest remaining cluster is 210 copies of a video falsely alleging that one's vote will not be counted if they press "Confirm vote" too quickly at the confirmation stage on the voting machine.

The next largest cluster is of 175 similar text messages. The messages in this cluster are on average 3,945 characters long and tell a "first-hand account" of alleged corruption and bribing of the TSE without which Bolsonaro would have supposedly won the election.

There are further clusters of videos making false allegations about discarded votes in Curitiba (134 videos) or fraudulent ballots being prepared by a political party in advance (131 videos).

Then 122 requests are about video where an influencer and self-titled leader of the New Aeon Church of Lucifer claims that different religions and entities linked to Satanism and Luciferianism had come together to guarantee the PT's victory in the first round of voting.

The next largest cluster is one of the 114 text messages alleging that "dead people voted for Lula." The average length of the messages in this cluster was 807 characters. After the initial allegation, the messages listed the names of various cities, their number of inhabitants, and the number of votes

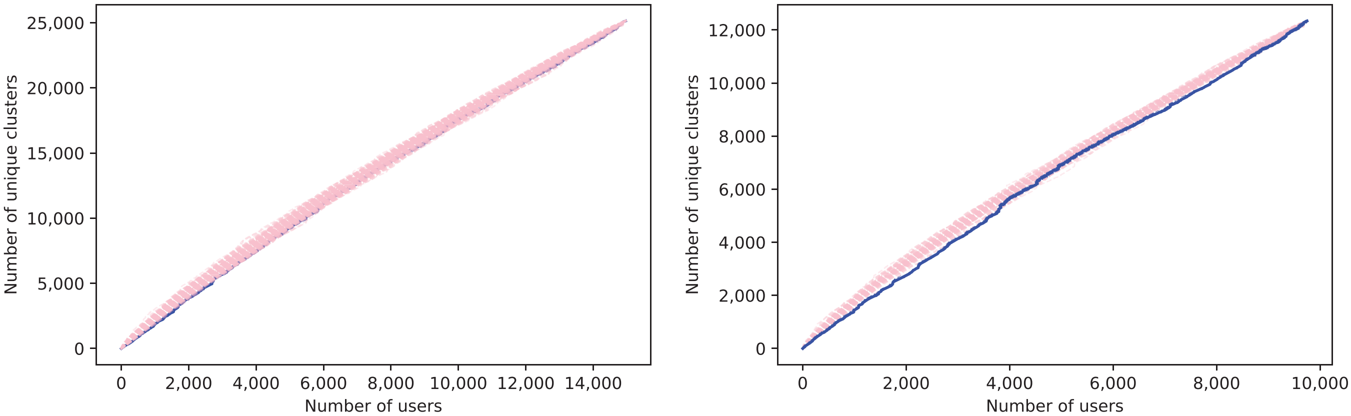

Figure 4. The relationship between the number of users and the amount of novel content (number of clusters) is mostly linear in the empirical data (solid blue line) and in random re-orders (pink dashed lines). The relationship is similar when considering all content (left) or only content that leads to a published fact-check (right).

cast for Lula. In each case, the number of votes cast for Lula was higher than the number of the inhabitants listed. The messages concluded by saying "there are 192 more Brazilian cities in which the dead resurrected to vote for Lula."

Other large clusters of *text* messages shared election polling predictions (94 messages), long lists of the good things Bolsonaro did for Brazil (77 messages), and "news" that protests had been successful in creating a new court that would oversee the Supreme Federal Court or STF (74 messages).

TSE Tipline

The data sent to the fact-checking feature of the TSE WhatsApp bot had far more text queries than the misinformation tiplines operated by fact-checkers. Between September 1, 2022, and November 15, 2022 inclusive the TSE fact-checking feature received 223,621 queries. The majority, 83%, are text messages, while videos account for 7%, images for 6%, and audios for 3%. This is partly explained by the fact that the feature only accepted text messages when it first launched. Multimedia messages were accepted from September 21, 2022.

The text messages sent to the TSE bot are also significantly shorter on average than those sent to the fact-checker tiplines. While the fact-checker tipline text messages have an average length of 637 characters (SD: 1,152), the text messages sent to the TSE bot have an average of 35 characters (SD: 57). In contrast to the long messages sent to fact-checkers, the TSE messages often reassemble search queries. The largest cluster is short variations of "ballot security" ("segurança das urnas") with 7,211 messages, but it must be noted that this was given as an example query by the bot, and some people clearly copied and pasted the example to try out the bot. In some instances the quotations marks around the phrase or the word "exemplo" [example] are included. Nevertheless, there are 204 variations of this query, suggesting that there was also genuine interest. There are also 956 searches for "fraude nas urnas" [ballot fraud]. The other top clusters include "[e-]título" ("voter id," 3,084 messages), "justificar voto" ("justifying vote," 1,997 messages), "local de votação" ("place of voting," 1,762), and "voto em trânsito" ("vote in transit," 1,572). Despite being told the fact-checking feature was exclusively for concerns about the integrity of the elections, there are messages about candidates as well: "Lula" was submitted 857 times and "Bolsonaro" 780 times. Variations of "Lula ladrão" [Lula thief] and "Lula é inocente?" [Is Lula innocent] were sent 2,858 times. After the elections, searches for "resultado" [result] increased with 751 queries overall.

Overall, overlap with the fact-checking misinformation tiplines is low. Of the 8,389 videos appearing across the tiplines and TSE bot, only 1,533 (18%) appear in both. Similarly, only 1% of images and less than 0.01% of text claims appear in both sources. For the items appearing in both the tiplines and the TSE bot, there is a weak, but positive and statistically significant correlation between the number of appearances in both (Figure 5). The correlation is .67 for videos, .55 for images, and .32 for text messages.

Kwai and WhatsApp

We identified 13,068 unique Kwai users. As in the WhatsApp data, the distribution of posts per user is heavy-tailed: 82% of the users appear in our dataset with just one post and 10% with two. One user had 223 videos in our dataset and ran a channel exclusively dedicated to the Brazilian elections.

We embed the video descriptions and manually inspect the largest clusters. We find 9,068 semantic clusters for the 15,017 video descriptions. The largest cluster contains 301 posts, and its keywords mirror the broad topic of our research ("Eleicoes2022" and variations). The second most important cluster identified has 85 posts and contains words associated with the Brazilian army ("exercito"). For comparison purposes, the top clusters associated with the words "Lula" and "Bolsonaro" had 55 and 54 posts, respectively.

To compare the video content present in the WhatsApp and Kwai datasets, we produced TMK embeddings for all 35,701 videos from both data sources and used the `tmk-clusterize` executable to cluster those embeddings, with the recommended .7 threshold for both level-1 and -2 similarity measures. We find that increasing the threshold to a more stringent value of .9 does not result in any relevant difference in our results, suggesting that the findings are robust to a range of threshold choices. In total, we find 20,887 clusters, of which 17,364 are singleton clusters, i.e., clusters with only one video. Of those 17,364 singleton clusters, 6,338 videos are from WhatsApp, and 11,026 videos are from Kwai (a 36%/64% split). The left panel of Figure 6 shows a histogram of cluster size for all clusters.

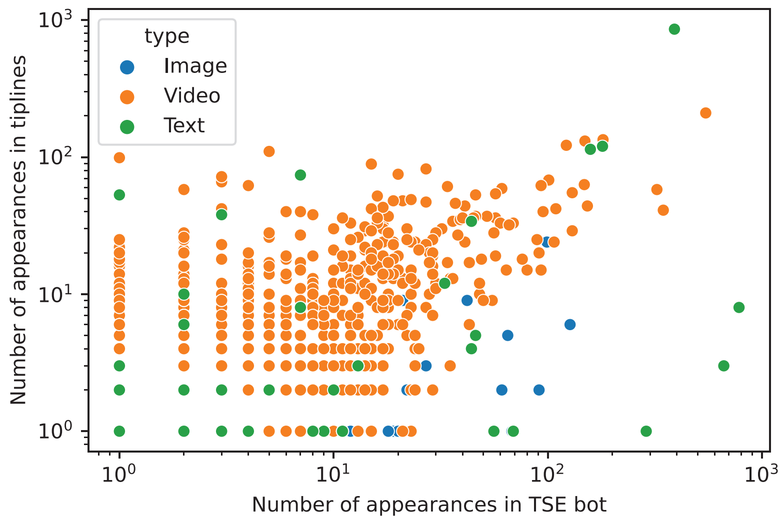

Figure 5. Overlap between content submitted to the tiplines and TSE is low, but there is a weak, positive correlation between the number of times a video, image, or text item is sent to the TSE bot and the misinformation tiplines.

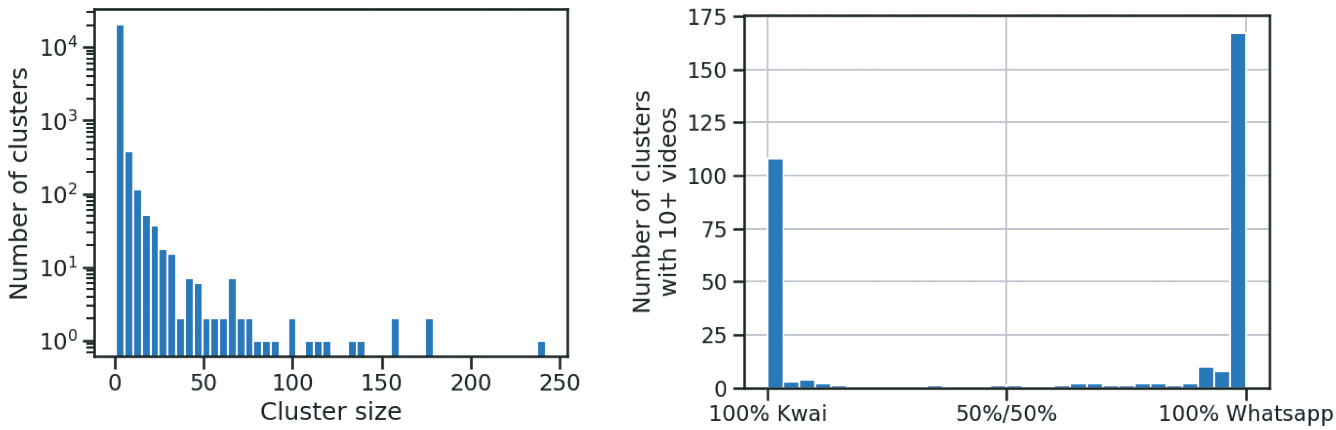

Figure 6. Left: Distribution of cluster sizes for the clusters of TMK embeddings of videos from WhatsApp and Kwai. Right: Histogram of the percentage of WhatsApp videos per cluster, for the 320 clusters with 10 or more videos.

For the non-singleton clusters, we find that the majority of clusters are composed of either nearly 100% of Kwai videos or nearly 100% of WhatsApp videos. This is illustrated on the right panel of Figure 6, which shows a histogram of the percentage of WhatsApp videos per cluster for the 320 clusters with 10 or more videos: The distribution is bimodal, with the two largest peaks corresponding to clusters that are composed of videos exclusive to Kwai (corresponding to 32.2% of clusters) or exclusive to WhatsApp (49.7% of clusters).

Twitter and WhatsApp

We find embeddings and existing claim-matching algorithms have low precision for claim-matching the long WhatsApp text messages with short Twitter posts. As a result, we use a combination of quantitative methods and manual analysis to investigate the prevalence of the top four WhatsApp text messages on Twitter. The top cluster of text messages on WhatsApp alleges corruption at the Ministry of Justice or the TSE. The messages in this cluster made use of phrases such as "se as pessoas soubessem o que aconteceu nos bastidores do TSE ficariam enojadas" [If people knew what happened behind the scenes at TSE they would be disgusted] or, "se as pessoas soubessem o que aconteceu no ministerio da justica ficariam enojadas" [If people knew what happened at the Ministry of Justice they would be disgusted.

The expression "If people knew X, they would be disgusted" is a long-standing *snowclone* (a phrasal template or, effectively, a meme) in Brazil.[8] It dates back to misinformation claims related to the 1998 World Cup and has been adapted by Brazilian Internet users in various contexts since then. Most tweets in this cluster were short, often repeating the sentence with little variation and no further detail—probably due to the meme status of the sentence. We examined the results returned by the semantic models and then filtered them to only those containing "enojadas" and "pessoas soubessem." With this filtering, we found 23 matching tweets. The oldest matching tweet in our data was published on October 3, 2022, while the first occurrence of the message in the WhatsApp tiplines

[8] The term "snowclone" was introduced by American linguists Geoffrey K. Pullum and Glen Whitman to describe textual templates and cliché frames (Pullum, 2004).

was on September 5. While the matching tweets often do not provide any real detail, the WhatsApp messages are longer and contain lists of claims about bribes, corruption, and other unsubstantiated assertions.

The second largest cluster of WhatsApp text messages is about “cidades que ate os mortos votaram em lula” [cities where even the dead voted for Lula]. After manual inspection, we filtered messages returned with the semantic models to only those containing “mortos” [dead] and “votar” [vote] and found 284 matching tweets. The earliest was published at 1:28 AM UTC on October 4, 2022, while the earliest submission on the WhatsApp tipline was sent a little less than an hour earlier at 00:48 AM UTC on the same day. While the first tweet on Twitter was labeled as “misleading” and could no longer be found when searching for “cidades que ate os mortos votaram em lula” on the Twitter website despite that exact phrase appearing in the tweet. Nonetheless, at the time of writing in January 2023, there were many very similar tweets with the same phrase that were not labeled as misleading (Figure 7). One tweet accessible via the search feature on Twitter includes screenshots of the messages on WhatsApp.

The third largest WhatsApp cluster was a message describing (incorrect) partial poll results in each Brazilian state, and asking other WhatsApp users to share the content with their friends and Bolsonaro supporters. It claims that mainstream news sources and social media platforms would not publicize the results otherwise. The matches in this cluster are mixed: The top matches in the cluster, i.e., the tweets with the highest similarity to the WhatsApp message, state support for Bolsonaro but do not mention any polls. Further matches, on the other hand, mention poll results but do not always state support for a candidate. This divide is a good illustration of the challenge in clustering embeddings of messages of variable length: Since many of the longer WhatsApp messages contain a header or a call to action (e.g., “everybody please read this” or “please share this message”), as well as the other contextual elements such as expressing support for a candidate, semantic similarity-induced clusters are likely to reflect these characteristics of the messages, rather than only producing topic-centered clusters.

Finally, the fourth largest cluster is centered around a WhatsApp message listing Bolsonaro’s alleged achievements as president. The full WhatsApp message lists 60 accomplishments and has a length of over 3,600 characters, which is more than 10 times the 280-character limit for a tweet. In this case, even though no single tweet accurately matches the whole list from the WhatsApp messages, a Twitter search for parts of the list returns multiple tweets containing the individual claims on the list, as well as Twitter threads containing multiple tweets, each with one or a few claims from the list.

Discussion

Our investigation into the 2022 Brazilian general elections reveals the multifaceted nature of misinformation across different platforms. While common claims exist across WhatsApp, Kwai, and Twitter, their presentation varies significantly consistent with affordance theory (Gibson, 1979).

Political engagement through WhatsApp in Brazil (Milan & Barbosa, 2020; Rossini, Baptista, Oliveira, & Stromer-Galley, 2021) and the spread of deceptive information during Brazilian elections (Machado, Kira, Narayanan, Kollanyi, & Howard, 2019; Resende et al., 2019) have been extensively documented. Results from a survey with Brazilian users show how both factors are intertwined: Participating in political groups on WhatsApp is correlated with belief in electoral misinformation (Rossini, Mont’Alverne, & Kalogeropoulos, 2023). Therefore, combating misinformation online, including social media and messaging apps, was the top priority for fact-checking organizations during the 2022 elections (Cazzamatta & Santos, 2023). The misinformation tiplines we study here run by the TSE and fact-checking organizations during the 2022 elections offer a privacy-preserving way to discover and respond to misinformation in closed messaging spaces. Unlike searching open social media, tiplines depend on citizen participation.

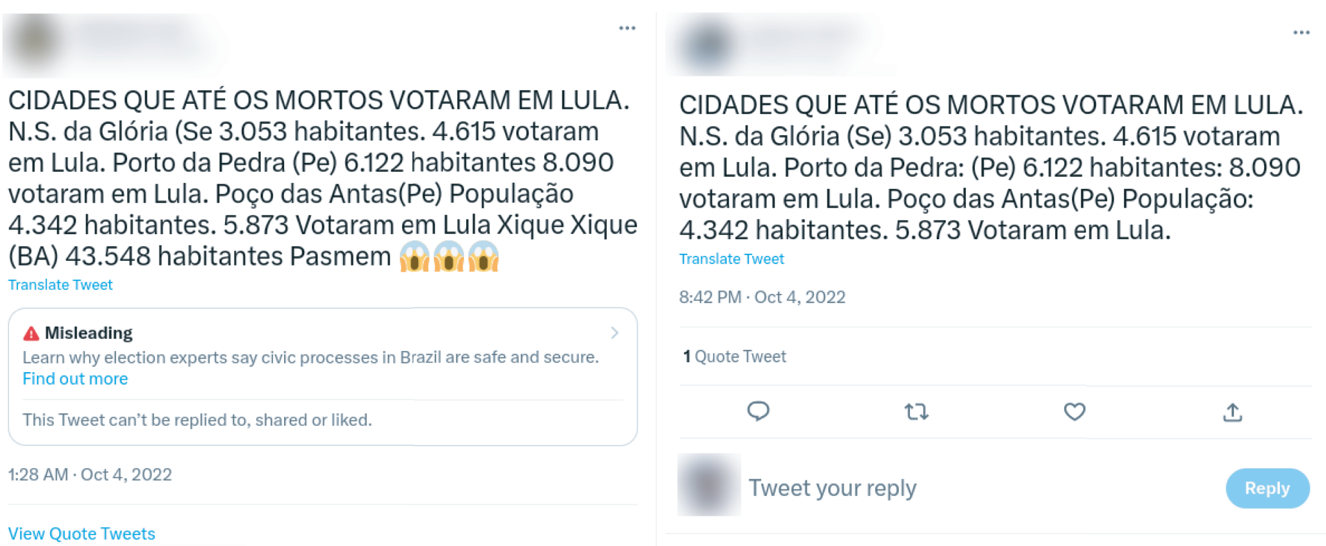

Figure 7. Tweets alleging there where “cities where even the dead voted for Lula.” While some tweets were labeled as misleading (left) and others were not (right).

Our comparative analysis of WhatsApp tiplines revealed that most users submit only one message, although a smaller number of "power-users" submit many messages. This distribution mirrors the heavy-tailed patterns found for content creation on other platforms (e.g., Panciera, Halfaker, & Terveen, 2009) and online political participation more generally (Margetts, John, Hale, & Yasseri, 2015). Of the users who submit multiple messages, most interact exclusively with one tipline, which provides initial evidence to suggest that each fact-checking organization serves a different audience. This further suggests that the misinformation sent to each tipline will be biased by the characteristics of each fact-checking organization's audience and not representative of misinformation from the whole political spectrum. We found that the amount of unique content is highly correlated with the number of users, suggesting that WhatsApp tiplines are not reaching a saturation point: as the number of users interacting with a tipline increases, so too does the amount of novel, unseen content.

While there are common claims across the platforms included in this study, there are also significant differences across platforms. Our finding of limited overlap in the videos on Kwai and WhatsApp might come as a surprise to WhatsApp researchers who anecdotally observe Kwai videos circulating on WhatsApp. However, it is important to remember that the WhatsApp dataset we analyze is not a sample of WhatsApp public groups, but rather a collection of crowd-sourced (or crowd-selected) content sent to fact-checking tiplines. Similarly, since Kwai does not offer a programmable interface for accessing public content, the Kwai videos analyzed might not be representative of all election-related content on the platform.

Our detailed analysis of specific misinformation clusters reveals the nuanced differences in content across platforms. The use of the "snowclone" phrasal template in WhatsApp messages and its brief appearances on Twitter exemplify how the same core message adapts to platform constraints and user behaviors. The longer format of WhatsApp allows for elaboration and multiple claims, whereas Twitter's character limit often reduces these messages to their most basic and meme-like form. This adaptation of content across platforms is a clear indication of how users tailor their communication based on their perceived audiences. This is influenced by platform-specific characteristics as well as by the groups users participate in and how they see themselves relating to them (i.e., their social identities), their emotional engagement, and group dynamics in general. Platforms with a timeline, such as Kwai and Twitter, may allow for a more fluid imagined audience, whereas closed-group platforms such as WhatsApp create multiple smaller audiences, allowing for multiple degrees of belonging and enacting of a social identity.

The study highlights the importance of considering platform-specific affordances and intrinsic differences in content presentation when analyzing the spread of misinformation claims across platforms. These differences create challenges for claim-matching algorithms (Kazemi, Garimella, Gaffney, & Hale, 2021; Shaar, Babulkov, Da San Martino, & Nakov, 2020) to match similar claims across different platforms. While we focused on individual messages as our unit of analysis, at times, one WhatsApp message might contain multiple allegations. Therefore, looking at posts or WhatsApp submissions as individual claims might not always be the right unit of analysis, as often the same political narratives will be spread through multiple versions of the same claims, each with small changes in form and content across platforms, all of which might be significantly different at an individual level but might still be pushing a common narrative. Additionally, the growth of video-based platforms, such as Kwai or TikTok, demands different analytical tools than text-based platforms. While TMK embeddings have proven useful in identifying identical videos across platforms, they do not address semantic elements of video content, i.e. when videos are on the same topic or frame an issue in the same way. TMK embeddings hash the visual channel of the videos and ignore audio, which is likely to provide important context. Multimodal claim-matching methods are an ongoing need for misinformation response.

Conclusion

The cross-platform movement of content, while evident in our data, does not lead to a homogeneous set of misinformation claims across platforms. Even when content does migrate from one platform to another, it is often reshaped or recontextualized to fit the new platform's affordances. Messages may change length or format when tailored to different audiences.

Our findings have important practical implications for fact-checkers and academics. Fact-checking organizations must adopt approaches to discovering and tracking misinformation claims in how they are presented to their audiences on multiple platforms. It is insufficient to monitor only the "easy-to-access" platforms and focus on the main misinformation claims there. Especially in Brazil where WhatsApp use is far greater than Twitter, it is necessary to develop approaches to identify misinformation claims directly on WhatsApp. Thankfully several options are developing: Fact-checking organizations can run tiplines like those used in this paper, scrape content from large, "public" groups (Garimella & Tyson, 2018), and work with data donation initiatives.[9] Overall, we hope that our research encourages further analysis of cross-platform misinformation, while also encouraging more research in cooperation with fact-checking organizations.

Acknowledgements

We are grateful to the fact-checking organizations for supporting this research and all the tireless work they did helping people locate high-quality trusted information. We are deeply grateful to our reviewers and editors for their helpful comments on earlier versions of our manuscript. Meedan receives funding from Meta and other technology companies, but the content of this manuscript was prepared with no input from these funders. S.A.H. and C.C. received funding from The Alan Turing Institute.

Biographical Note

Dr. Scott A. Hale (DPhil) is the Director of Research at Meedan, a global technology non-profit building tools for

[9] Example data donation initiatives include research phone apps (e.g., WhatsViral: https://play.google.com/store/apps/details?id=com.rutgers.whatsviral), Mozilla's YouTube Regrets (https://foundation.mozilla.org/en/youtube/findings/), tracking.exposed (https://tracking.exposed/), and the National Internet Observatory (Meyer, Basl, Choffnes, Wilson, & Lazer, 2023).

fact-checking and media annotation. He is also an Associate Professor and a Senior Research Fellow at the Oxford Internet Institute, University of Oxford, where he develops and applies NLP methods to increase equitable access to quality information.

Adriano Belisario (MSc) studied the MSc in Social Data Science course at the Oxford Internet Institute, University of Oxford. He also has collaborated on research projects with the Medialab of the Federal University of Rio de Janeiro and taught data journalism postgraduate courses in Brazil.

Ahmed Nasser Mostafa (MSc) is a Machine Learning Engineer at Meedan.

Dr. Chico Camargo (DPhil) is a Lecturer in Computer Science at the University of Exeter, a research associate at the Oxford Internet Institute, University of Oxford, and Computational Social Science Theme Lead at the Exeter Institute for Data Science and Artificial Intelligence. He is also a visiting professor at the College of Liberal Arts, Ewha Womans University, Seoul, South Korea.

References

Altay, S., Berriche, M., & Acerbi, A. (2023). Misinformation on misinformation: Conceptual and methodological challenges. *Social Media + Society*, *9*(1). doi:10.1177/20563051221150412

Anita Baptista, E., Rossini, P., Veiga de Oliveira, V., & Stromer-Galley, J. (2019). A circulação da (des) informação política no WhatsApp e no Facebook. *Lumina*, *13*(3), 29–46. doi:10.34019/1981-4070.2019.v13.28667

Asian American Disinformation Table. (2022). Power, platforms, politics: Asian Americans and disinformation landscape report. Retrieved from https://www.asianamdisinfo.org/wp-content/uploads/2022/08/AsianAmDisinformation_LandscapeReport2022.pdf

Blank, G., Dutton, W. H., & Lefkowitz, J. (2020). OxIS 2019: The rise of mobile internet use in Britain. Retrieved from https://ssrn.com/abstract=3538301

Brummette, J., DiStaso, M., Vafeiadis, M., & Messner, M. (2018). Read all about it: The politicization of "fake news" on Twitter. *Journalism & Mass Communication Quarterly*, *95*(2), 497–517. doi:10.1177/1077699018769906

Cazzamatta, R., & Santos, A. (2023). Checking verifications during the 2022 Brazilian run-off election: How fact-checking organizations exposed falsehoods and contributed to the accuracy of the public debate. *Journalism*. doi:10.1177/14648849231196080

Cetic.br. (2021). *Executive summary: ICT households survey 2021*. Retrieved from https://cetic.br/media/docs/publicacoes/2/20221121130127/executive_summary_ict_households_2021.pdf

Cihon, P., & Yasseri, T. (2016). A biased review of biases in Twitter studies on political collective action. *Frontiers in Physics*, *4*(34), 91–98. doi:10.3389/fphy.2016.00034

Deck, A., & Marasciulo, M. (2022). TikTok's biggest Chinese competitor bets big on Brazil. *Rest of World*. Retrieved from https://restofworld.org/2022/tiktok-competitor-kwai-brazil/

Dwoskin, E. (2023). Come to the 'war cry party': How social media helped drive mayhem in Brazil. *Washington Post*. Retrieved from https://www.washingtonpost.com/technology/2023/01/08/brazil-bolsanaro-twitter-facebook/

Garimella, K., & Tyson, G. (2018). *WhatApp doc? A first look at WhatsApp public group data*. Proceedings of the International AAAI Conference on Web and Social Media, *12*(1), 511–517. doi:10.1609/icwsm.v12i1.14989

Gava, M. (2022). Estudo reúne dados sobre o uso de redes sociais no Brasil. Retrieved from https://www.capterra.com.br/blog/3007/uso-redes-sociais

Gibson, J. J. (1979). *The ecological approach to visual perception*. Houghton Mifflin.

Ha, L., Andreu Perez, L., & Ray, R. (2021). Mapping recent development in scholarship on fake news and misinformation, 2008 to 2017: Disciplinary contribution, topics, and impact. *American Behavioral Scientist*, *65*(2), 290–315. doi:10.1177/0002764219869402

Hernon, P. (1995). Disinformation and misinformation through the internet: Findings of an exploratory study. *Government Information Quarterly*, *12*(2), 133–139. doi:10.1016/0740-624X(95)90052-7

Jerit, J., & Zhao, Y. (2020). Political misinformation. *Annual Review of Political Science*, *23*(1), 77–94. doi:10.1146/annurev-polisci-050718-032814

Johansson, P., Enock, F., Hale, S., Vidgen, B., Bereskin, C., Margetts, H., & Bright, J. (2022). How can we combat online misinformation? A systematic overview of current interventions and their efficacy. *arXiv*. https://arxiv.org/abs/2212.11864, preprint: not peer reviewed.

Jungherr, A. (2018). *Normalizing digital trace data*. New York: Routledge.

Kazemi, A., Garimella, K., Shahi, G. K., Gaffney, D., & Hale, S. A. (2022). Tiplines to uncover misinformation on encrypted platforms: A case study of the 2019 Indian general election on WhatsApp. *Harvard Kennedy School (HKS) Misinformation Review*, *3*(1), 1–17. doi:10.37016/mr-2020-91

Kazemi, A., Garimella, K., Gaffney, D., & Hale, S. (2021). *Claim matching beyond English to scale global fact-checking*. Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Vol. 1: Long Papers), 4504–4517. doi:10.18653/v1/2021.acl-long.347

Kemp, S. (2022). *Digital 2022: Brazil*. Data reportal. Retrieved from https://datareportal.com/reports/digital-2022-brazil

Lamb, R., King, J. L., & Kling, R. (2003). Informational environments: Organizational contexts of online information use. *Journal of the American Society for Information Science and Technology*, *54*(2), 97–114. doi:10.1002/asi.10182

Lelo, T. (2022). The rise of the Brazilian fact-checking movement: Between economic sustainability and editorial independence. *Journalism Studies*, *23*(9), 1077–1095. doi:10.1080/1461670X.2022.2069588

Litt, E. (2012). Knock, knock. who's there? The imagined audience. *Journal of Broadcasting & Electronic Media*, *56*(3), 330–345. doi:10.1080/08838151.2012.705195

Machado, C., Kira, B., Narayanan, V., Kollanyi, B., & Howard, P. (2019). *A study of misinformation in WhatsApp groups with a focus on the Brazilian presidential elections*. Companion Proceedings of The 2019 World Wide Web Conference, 1013–1019. doi:10.1145/3308560.3316738

Margetts, H., John, P., Hale, S., & Yasseri, T. (2015). *Political turbulence: How social media shape collective action*. Princeton: Princeton University Press.

Meyer, M. N., Basl, J., Choffnes, D., Wilson, C., & Lazer, D. M. J. (2023). Enhancing the ethics of user-sourced online data collection and sharing. *Nature Computational Science*, *3*(8), 660–664. doi:10.1038/s43588-023-00490-7

Milan, S., & Barbosa, S. (2020). Enter the WhatsApper: Reinventing digital activism at the time of chat apps. *First Monday*. doi:10.5210/fm.v25i12.10414

Newman, N., Fletcher, R., Robertson, C. T., Eddy, K., & Nielsen, R. K. (2022). *Digital news report 2022*. Reuters Institute, University of Oxford. Retrieved from https://reutersinstitute.politics.ox.ac.uk/digital-news-report/2022

Norman, D. (2013). *The design of everyday things: Revised and expanded edition*. Basic Books.

Osmundsen, M., Bor, A., Vahlstrup, P. B., Bechmann, A., & Petersen, M. B. (2021). Partisan polarization is the primary psychological motivation behind political fake news sharing on Twitter. *American Political Science Review*, *115*(3), 999–1015. doi:10.1017/S0003055421000290

Paiva, F. (2020). WhatsApp alcança presença recorde em 99% dos smartphones no Brasil. *Mobile Time*. Retrieved from https://www.mobiletime.com.br/noticias/27/02/2020/whatsapp-alcanca-presenca-recorde-em-99-dos-smartphones-no-brasil/

Paiva, F. (2022). Panorama: Uso de apps no Brasil. *Mobile Time*. Retrieved from https://www.mobiletime.com.br/pesquisas/download/35312/

Panciera, K., Halfaker, A., & Terveen, L. (2009). *Wikipedians are born, not made: A study of power editors on Wikipedia*. Proceedings of the 2009 ACM International Conference on Supporting Group Work. Association for Computing Machinery, 51–60. doi:10.1145/1531674.1531682

Pantazi, M., Hale, S., & Klein, O. (2021). Social and cognitive aspects of the vulnerability to political misinformation. *Political Psychology*, *42*(S1), 267–304. doi:10.1111/pops.12797

Pennycook, G., Epstein, Z., Mosleh, M., Arechar, A. A., Eckles, D., & Rand, D. G. (2021). Shifting attention to accuracy can reduce misinformation online. *Nature*, *592*(7855), 590–595. doi:10.1038/s41586-021-03344-2

Pereira, A., Harris, E., & Van Bavel, J. J. (2023). Identity concerns drive belief: The impact of partisan identity on the belief and dissemination of true and false news. *Group Processes & Intergroup Relations*, *26*(1), 24–47. doi:10.1177/13684302211030004

Pullum, G. K. (2004). Snowclones: Lexicographical dating to the second. *Language Log*. Retrieved from http://itre.cis.upenn.edu/˜myl/languagelog/archives/000350.html

Rathje, S., Van Bavel, J. J., & Van Der Linden, S. (2021). Out-group animosity drives engagement on social media. *Proceedings of the National Academy of Sciences of USA*, *118*(26), e2024292118. doi:10.1073/pnas.2024292118

Recuero, R. (2020). #FraudenasUrnas: Estratégias discursivas de desinformação no Twitter nas eleições 2018. *Revista Brasileira de Linguística Aplicada*, *20*, 383–406. doi:10.1590/1984-6398202014635

Reimers, N., & Gurevych, I. (2019). *Sentence-BERT: Sentence embeddings using Siamese BERT-networks*. Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), Association for Computational Linguistics, 3982–3992. doi:10.18653/v1/D19-1410

Reis, J., Melo, P., Garimella, K., & Benevenuto, F. (2020). Can WhatsApp benefit from debunked fact-checked stories to reduce misinformation? *Harvard Kennedy School (HKS) Misinformation Review*. doi:10.37016/mr-2020-035

Reis, J. C. S., Melo, P., Garimella, K., Almeida, J. M., Eckles, D., & Benevenuto, F. (2020). A dataset of fact-checked images shared on WhatsApp during the Brazilian and Indian elections. *Proceedings of the International AAAI Conference on Web and Social Media*, *14*, 903–908. doi:10.1609/icwsm.v14i1.7356

Resende, G., Melo, P., Sousa, H., Messias, J., Vasconcelos, M., Almeida, J., & Benevenuto, F. (2019). *(Mis)information dissemination in WhatsApp: Gathering, analyzing and countermeasures*. The World Wide Web Conference, Association for Computing Machinery, USA, 818–828. doi:10.1145/3308558.3313688

Rossini, P., Baptista, É. A., Oliveira, V. V. d., & Stromer-Galley, J. (2021). Digital media landscape in Brazil: Political (mis)information and participation on Facebook and WhatsApp. *Journal of Quantitative Description: Digital Media*, *1*(2021), 1–27. doi:10.51685/jqd.2021.015

Rossini, P., Mont'Alverne, C., & 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), 1–16. doi:10.37016/mr-2020-115

Ruths, D., & Pfeffer, J. (2014). Social media for large studies of behavior. *Science*, *346*(6213), 1063–1064. doi:10.1126/science.346.6213.1063

Shaar, S., Babulkov, N., Da San Martino, G., & Nakov, P. (2020). *That is a known lie: Detecting previously fact-checked claims*. Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, Association for Computational Linguistics, 3607–3618. doi:10.18653/v1/2020.acl-main.332

Starbird, K., Arif, A., Wilson, T., Van Koevering, K., Yefimova, K., & Scarnecchia, D. (2018). *Ecosystem or echo-system? Exploring content sharing across alternative media domains*. Proceedings of the International AAAI Conference on Web and Social Media, *12*(1), Association for the Advancement of Artificial Intelligence, USA. doi:10.1609/icwsm.v12i1.15009

Tarouco, G. (2023). Brazilian 2022 general elections: Process, results, and implications. *Revista Uruguaya de Ciencia Política*, *32*(1), 154–168. doi:10.26851/rucp.32.1.7

Thorson, K., Driscoll, K., Ekdale, B., Edgerly, S., Thompson, L. G., Schrock, A., Swartz, L., Vraga, E. K., & Wells, C. (2013). YouTube, Twitter and the Occupy movement: Connecting content and circulation practices. *Information, Communication & Society*, *16*(3), 421–451. doi:10.1080/1369118X.2012.756051

TSE. (2022). Eleições 2022: TSE e WhatsApp discutem medidas para enfrentamento da desinformação. Retrieved from https://www.tse.jus.br/comunicacao/noticias/2022/Janeiro/eleicoes-2022-tse-e-whatsapp-discutem-medidas-para-enfrentamento-da-desinformacao

Van Bavel, J. J., Rathje, S., Vlasceanu, M., & Pretus, C. (2024). Updating the identity-based model of belief: From false belief to the spread of misinformation. *Current Opinion in Psychology*, *56*(101787). doi:10.1016/j.copsyc.2023.101787

Wilson, T., & Starbird, K. (2020). Cross-platform disinformation campaigns: Lessons learned and next steps. *Harvard Kennedy School Misinformation Review*, *1*(1), 1–11. doi:10.37016/mr-2020-002