Special Issue: Trust and Safety on Social Media

SM+S
social media + society

Social Media + Society
July-September 2023: 1–13
© The Author(s) 2023
Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/20563051231196866
journals.sagepub.com/home/sms
Sage

A Workflow to Detect, Monitor, and Update Lists of Coordinated Social Media Accounts Across Time: The Case of the 2022 Italian Election

Fabio Giglietto, Giada Marino, Roberto Mincigrucci, and Anna Stanziano

Abstract

Information operations that target public opinion often exploit breaking news, crises, and elections by using coordinated social media actors to disseminate problematic content. These events often reveal the relationships between actors, prompting the creation of lists of malicious actors and news sources. However, relying on outdated lists may underestimate the prevalence and impact of such operations. This article presents a novel workflow to detect, monitor, and update lists of coordinated social media actors during and beyond peak activity periods. Using this approach, known problematic actors are constantly monitored, allowing the detection of new actors and the update of the monitored pool. The workflow was applied to the 2022 Italian snap election, leveraging previous research on coordinated inauthentic behavior during the 2018 and 2019 Italian elections. The initial list of 435 coordinated accounts was monitored, surfacing 1,022 overly shared or commented political posts, 272 coordinated links, and detecting 66 political and 554 generic coordinated accounts not previously listed. Three case studies were identified: one politically motivated, one click-economy driven, and one religiously motivated operation. The article discusses the implications of this approach, its limitations, and potential future work.

Keywords

information operations, coordinated inauthentic behavior, social media monitoring, problematic content detection, digital political communication

Introduction

Actors seeking to manipulate public opinion are drawn to the attention generated by breaking news, crises, and elections. During these periods, the intensity of their activities increases, following consolidated tactics (Donovan et al., 2022; Tripodi, 2022), techniques (Golebiewski & Boyd, 2018), and procedures (Dawson & Innes, 2019). These patterns make information operations during such times easier to spot and the relationships between different actors more discernible. Consequently, most lists of known assets (social media accounts, internet domains, etc.) employed by these operations are collected during these events.

The compilation of these lists depends on the labor-intensive efforts of fact-checkers and debunkers. List-based approaches assume that actors who have been repeatedly caught sharing false, misleading, hyperpartisan, and manipulated content—problematic information in the broader sense proposed by Caroline Jack (2017) and adopted throughout this article—will persist in doing so in the future, thereby rendering them problematic as a whole. Once compiled, other studies extensively reuse these lists (Allcott & Gentzkow, 2017; Guess et al., 2019). However, as information operations constantly attempt to avoid detection and may change their social media assets due to suspension or deletion, studies based on outdated lists may underestimate the prevalence and impact of the observed phenomenon (Yang et al., 2023). This article introduces a novel—content

University of Urbino, Italy

Corresponding Author:
Fabio Giglietto, Department of Communication Sciences, Humanities and International Studies (DISCUI), University of Urbino, Via A. Saffi, 15, Urbino (PU) 61029, Italy.
Email: fabio.giglietto@uniurb.it

and actor agnostic, targeting behaviors rather than individuals or specific content—workflow devised to detect, monitor, and update lists of social media actors during and beyond the peak of their activities.

The workflow specifically targets coordinated behavior. Coordination, in and of itself, cannot be deemed categorically problematic (Magelinski et al., 2022). However, an increasing number of studies have linked coordination on social media to online influence operations (Hristakieva et al., 2022; Starbird et al., 2019).

Posts created by accounts performing this behavior during a specific period (e.g., a political campaign) are constantly monitored, and best-performing posts are used to identify additional actors coordinating to spread the same content. In other words, the behavior of known coordinated actors is utilized to detect new and/or additional social media assets. The newly discovered actors are then added to the pool of monitored accounts.

We applied this workflow to study the 2022 Italian snap election. Building on the results of previous works that examined coordinated inauthentic behavior during the 2018 and 2019 Italian elections (Giglietto et al., 2020a; Giglietto, Righetti, & Marino, 2019), we compiled an initial list of 435 known coordinated accounts (238 Facebook Pages, 196 Facebook public groups, and 1 Instagram account). Following the posts published by these accounts every 6 hr from July 28 to September 25, 2022 (election day), we identified 1,022 highly shared or commented political posts and 272 coordinated links in near-real-time. In addition, we detected 66 coordinated political accounts (20 Facebook Pages and 46 public groups) and 554 generic (political and non-political) coordinated accounts (406 Facebook Pages and 148 public groups) not listed in our initial set of accounts.

In the following section, we contextualize the motivation for this work within the existing scientific literature by surveying recent studies that employed a list-based approach to examine the prevalence and spread of problematic information. The "Method" section details the workflow and its implementation deployed to study the 2022 Italian election. Finally, we describe three information operations (one politically motivated, one click-economy driven, and the last religiously motivated) identified by analyzing the content and accounts that surfaced.

Literature Review

The Limits of the Content-Based Approaches

Detecting, tracking, and measuring the circulation of problematic information on social media is essential for understanding its prevalence and ultimately finding effective strategies to minimize its detrimental effects on democracy (Benkler, 2019). Despite extensive interdisciplinary efforts (Righetti et al., 2022), several challenges hinder attempts to provide compelling and undisputed findings that policymakers can use to shape public and private policies.

Three main factors limit content-based approaches to detection: the phenomenon's complexity, speed, and scale. The intrinsic complexity and multifaceted nature of problematic information (Gleicher, 2018; Nimmo & Hutchins, 2023) obstruct systematic definitions and impede the development of effective automated detection methods. Problematic information extends beyond false content, as demonstrated by weaponized quotes (Marino & Giglietto, 2023) and factually accurate stories circulated to mislead the public about the prevalence of certain events (e.g., a collateral effect caused by a vaccination). Various attempts have been made to classify problematic information formats, particularly after 2016. Jack (2017) proposed a taxonomy that described misinformation, disinformation, propaganda, and gaslighting as "problematic information." Wardle and colleagues (2018) introduced a taxonomy of "information disorder" that encompasses various forms of problematic news content, and Donovan and colleagues (2020) provided an online resource with a repository of "media manipulation" cases, methodological tools, and definitions.

The detection speed is another challenge, as problematic content rapidly gains reach and engagement on social media. Zuckerberg and his team (2021) acknowledged that inaccurate, borderline, or harmful Facebook content generates higher engagement levels. Moreover, recent scientific reports (Integrity Institute, 2022; Matatov et al., 2022) indicate that inaccurate and harmful content typically receives the majority of its engagement within the first 24 hr of publication. This highlights the importance of prompt detection for mitigating the impact of misinformation and disinformation.

The final challenge is the issue of scale. Creating content is much easier than verifying, resulting in a structurally unbalanced competition between problematic content creators and fact-checkers. The popularization of generative artificial intelligence (AI) will further exacerbate this issue.

Given these challenges, most recent studies leverage the work on content by scaling it up at the level of actors/sources of problematic information.

Lists of Problematic News Sources

A common strategy to address the challenges in tracking problematic content involves creating lists of problematic social media actors or news sources (Allcott et al., 2019; Freelon et al., 2018; Guess et al., 2018; Houidi et al., 2019). Compiled by debunking organizations, journalists, and fact-checkers, these lists generally include websites that publish and spread problematic information, social media accounts, or groups (Forrester et al., 2019; Giglietto et al., 2020b). Lists are assembled automatically by incorporating repeated offenders (sources/actors that repeatedly publish or recirculate content deemed false by fact-checkers) or through manual investigation of the actor/source (NewsGuard, 2020).

Although this approach reduces the challenges to detection posed by speed and scale, it remains resource-intensive in terms of time, staff, and associated costs, as it requires journalists and fact-checkers to manually check news source content, investigate site ownership, and report their ratings. Moreover, malicious sources often vanish and reappear rapidly, adapting to the ever-changing online media environment (Bastos & Mercea, 2019). This results in frequently outdated lists and may lead to underestimating the phenomenon (Yang et al., 2023).

Computational Detection of Problematic Actors

A second and distinct approach involves computational methods to detect problematic actors. Automated accounts mimicking real users have increasingly become the focus of detection efforts, particularly on social media platforms (Gleicher, 2018; Santia et al., 2019). Malicious actors' strategies are not limited to using completely fake accounts; they may also use real user accounts while concealing their identities or intentions to propagate their ideas (Mazza et al., 2022).

Previous studies have demonstrated the activities of these networks of accounts, particularly engaged on divisive issues, such as the case described by Daniels (2009) of anti-abortion sites masked under the pro-choice tag or the false Islamist Facebook pages spreading anti-Muslim content analyzed by Farkas and colleagues (2018). Computational methods are used to detect social bots by examining account information or their network of ties (Cresci et al., 2016; Davis et al., 2016; Wu et al., 2020).

However, this branch of study has limitations. Using computational methods to detect malicious actors requires a narrowed conceptualization of what constitutes a malicious actor to reduce false positives (Giglietto et al., 2020b). In addition, tactics to hide fake accounts and bots have become increasingly sophisticated, making detection highly complex without extensive human supervision combined with machine learning.

Coordinated Inauthentic Behavior

In addition to authenticity, coordination is another feature considered in efforts to mitigate problematic information circulation. Organized activities on digital media are among the most recognizable aspects of user activity on the internet. Henry Jenkins' preliminary studies on fandom showed that online participation was significantly driven by user coordination of activities to achieve specific tasks, such as creating fan theories or finding reliable spoilers for their favorite audiovisual content (Jenkins, 2008). These activities fell into the definition of participatory culture.

For a long time, scholars and journalists considered this characteristic of online spaces positively disruptive. This was partly because civil rights activists and minority groups widely used it to organize activities, gain visibility, and promote their issues (Freelon & Wells, 2020). Despite this initial wave of optimism, years later, the same platforms and similar coordination techniques were employed by various malicious actors with diverse motivations (Cinelli et al., 2022; Starbird et al., 2019).

Balancing the protection of user voices and addressing malicious users at scale is considered an open challenge. This scenario may have led Meta to associate coordination with the concept of inauthenticity—misrepresentation of goal or identity—in the definition employed by their initiatives aimed at disrupting adversarial threats.

Research Questions

In designing this study, we considered the complexity of identifying the aforementioned problematic actors (malicious news sources) and behaviors (coordinated campaigns aimed at spreading problematic information), as well as the potential, challenges, and limitations of each detection method. With these factors in mind, this study aims to propose an approach for uncovering problematic information by detecting, tracking, and updating lists of coordinated social media actors. This approach focuses on the concept of Coordinated Link Sharing Behavior (CLSB) (Giglietto et al., 2020b; Gruzd et al., 2022) and utilizes computational methods to create and automatically update lists of Facebook entities that disseminate problematic information online.

We aim to evaluate this approach in the context of the recent Italian general elections. Therefore, we posed the following research questions:

RQ1: What types of coordinated behavior can this approach reveal in relation to the 2022 Italian general elections electoral campaign?

RQ2: Which actors were at the forefront of these coordinated campaigns?

Method

Process Overview

The workflow (Figure 1) requires one or more initial lists of known problematic social media accounts to be monitored. These lists can either be found in the existing literature or compiled ad hoc. The posts published by these accounts can be periodically monitored through a scheduled process by accessing the Application Programming Interface (APIs) of social media and social media analytics platforms. This process collects and evaluates the early performance of the content (actual) based on the historical performance of the post published by the accounts under examination (expected). The logic here is mutated from CrowdTangle's overperforming score (2016). The content of the overperforming posts

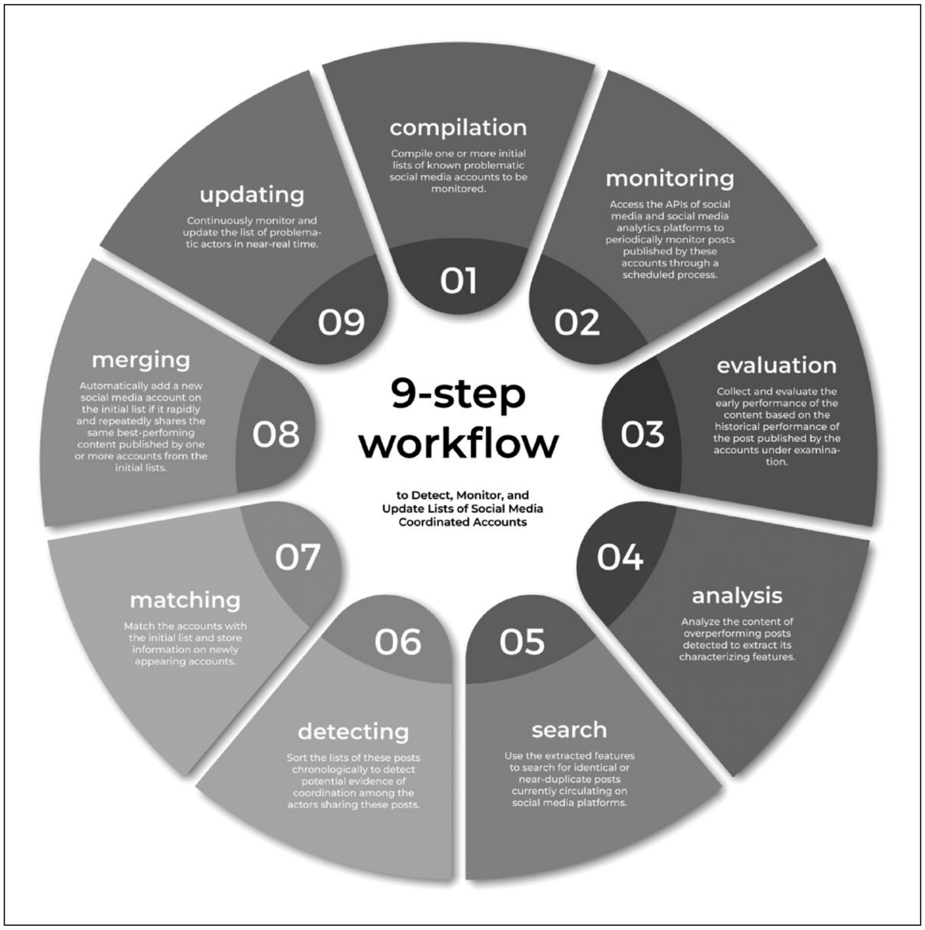

9-step workflow

to Detect, Monitor, and Update Lists of Social Media Coordinated Accounts

01 compilation

Compile one or more initial lists of known problematic social media accounts to be monitored.

02 monitoring

Access the APIs of social media and social media analytics platforms to periodically monitor posts published by these accounts through a scheduled process.

03 evaluation

Collect and evaluate the early performance of the content based on the historical performance of the post published by the accounts under examination.

04 analysis

Analyze the content of overperforming posts detected to extract its characterizing features.

05 search

Use the extracted features to search for identical or near-duplicate posts currently circulating on social media platforms.

06 detecting

Sort the lists of these posts chronologically to detect potential evidence of coordination among the actors sharing these posts.

07 matching

Match the accounts with the initial list and store information on newly appearing accounts.

08 merging

Automatically add a new social media account on the initial list if it rapidly and repeatedly shares the same best-perfoming content published by one or more accounts from the initial lists.

09 updating

Continuously monitor and update the list of problematic actors in near-real time.

Figure 1. A visualization of the circular process workflow.

detected is analyzed to extract its characterizing features (e.g., the text of the post, the text in the image of the post, and the link shared by the post). These features are then used to search for identical or near-duplicate posts currently circulating on social media platforms, including but not limited to the monitored accounts. Following the detection methods designed and implemented in tools such as CooRnet (Giglietto, Righetti, & Rossi, 2020), Coordination Network Toolkit (Graham & QUT Digital Observatory, 2020), and CooRTweet (Righetti & Balluff, 2023), the lists of these posts are sorted chronologically to detect potential evidence of coordination (e.g., identical or near-duplicate posts synchronously shared by multiple social media accounts) among the actors sharing these posts. When such behavior is detected, accounts are matched with the initial list, and information on newly appearing accounts is stored. When a new social media account surfaces multiple times (in a single or subsequent scheduled monitoring process), the account is automatically added to the initial list and thus begins to be constantly monitored. In other words, an account is added to the list of monitored accounts if it rapidly and repeatedly shares the same best-performing content published by one or more accounts from the initial lists.

This process detects new social media accounts added to known coordinated networks and contributes to keeping lists of problematic actors updated in near-real time.

Workflow Implementation for the 2022 Italian Election

About the Election. The 2022 Italian general election was called early on July 21, 2022, following the unexpected collapse of Mario Draghi's government. The election day was scheduled to be held on September 25, 2022.

Glossary

- CrowdTangle (Fraser, 2020) is a Meta-owned tool that tracks interactions on public content from Facebook Pages and groups, verified profiles, Instagram accounts, and subreddits. It does not include paid ads unless those ads began as organic, non-paid posts that were subsequently "boosted" using Facebook's advertising tools. It also does not include activity on private accounts or posts made visible only to specific groups of followers;
- CLSB is a specific coordinated activity performed by a network of accounts (Facebook Pages, public groups,

and verified public profiles or Instagram accounts) that repeatedly share (repetition threshold) the same news articles in a very short time (coordination interval) from each other;

- CooRnet (Giglietto, Righetti, & Rossi, 2020) is an R package that, given a set of URLs, detects networks of CLSB accounts on Facebook and Instagram. To do so, CooRnet analyzes data provided by CrowdTangle. CooRnet automatically estimated the coordination interval based on the dataset under examination. To put the researcher in charge, CooRnet optionally allows the user to manually specify his own coordination interval and repetition threshold (as a percentile of the weight distribution on the coordinated account graph);
- A CooRnet iteration (1) collects posts published in a certain period of time by a set of previously detected CLSB accounts, (2) extracts URLs shared by these posts, and (3) initiates a new CooRnet cycle using these sets of URLs. Each iteration surfaces new problematic content and detects additional coordinated accounts, thus refreshing the list of tracked accounts;
- Coordinated Image Text Sharing Behavior (CITSB) is a specific coordinated activity performed by a network of accounts (Facebook Pages, public groups, and verified public profiles or Instagram accounts) that repeatedly share images that contain the same text (e.g., image macros, memes) in a very short time from each other;
- Coordinated Message Sharing Behavior (CMSB) is an additional specific coordinated activity performed by a network of accounts (Facebook Pages, public groups, and verified public profiles or Instagram accounts) that repeatedly share posts that contain identical or near-duplicate (cosine similarity > .7) text in a very short time from each other.

Setup. To identify the initial list of accounts to be monitored, we leveraged previous works on coordinated link sharing networks during Italian elections (Giglietto et al., 2020a, 2020b; Giglietto, Righetti, & Marino, 2019). More specifically, we run a CooRnet iteration on links posted by three CrowdTangle lists of coordinated Facebook Pages and public groups detected, respectively, during the 2018 Italian general election, the 2019 Italian election for the European Union (EU) parliament (Giglietto, Righetti, & Marino, 2019), the post-electoral period between June and November 2019 (Giglietto et al., 2020b), and the acute phase of the COVID-19 crisis in Italy between January and October 2020 (Giglietto et al., 2022). To consider links shared in the first comment of the post—a known adaptation strategy where links will not be automatically detected by CrowdTangle and CooRnet (Giglietto, Terenzi, et al., 2020)—we additionally detected a list of CITSB accounts on July 28, 2022.

We then collected all the political posts[1] (N = 398,385) published by these accounts during the past 6 months before the day the election was called (January 23, 2021—July 22, 2022). Using CooRnet on these links (N = 73,842), we detected—using the coordination interval of 30 s estimated by the tool and a repetition threshold of 26 or more (0.995 percentile edge weight) rapid shares—435 accounts (238 Facebook Pages, 196 Facebook public groups, and 1 Instagram account).

The posts published by these accounts have been monitored every 6 hr (3 a.m., 9 a.m., 3 p.m., and 9 p.m.) by an R script scheduled with cronR (Wijffels, 2022) from July 28 to September 25, 2022. At each run, the script queries CrowdTangle posts/search API to retrieve a list of the top 100 political or unfiltered posts (all posts) sorted by CrowdTangle overperforming score (CrowdTangle Team, 2016). An additional list of 100 best-performing political/unfiltered posts published by the 10% of most frequently detected new coordinated accounts identified during the previous runs are also added to the monitored posts. The two lists of posts are then bound and deduplicated.

Links are then extracted from the collected posts and used to run a real-time CooRnet iteration. Coordinated links detected are stored in an online spreadsheet. The image text (a field provided by CrowdTangle by extracting text from image-type posts) and text of the post (message field) are additionally used to detect non-link-based forms of coordination (sharing images with the same text—CITSB—and sharing posts with the near-duplicate text—CMSB). Newly discovered coordinated accounts (CLSB, CITSB, and CMSB) are bound, deduplicated, and stored with their detection date.

The script calculates two additional metrics for each post: comment.shares.ratio (Giglietto, Valeriani, et al., 2019) and combined.metric (a synthetic measure of performance and comments/share balance obtained by multiplying the overperforming score by the comment.share.ratio). Posts are then sorted by combined.metric. The top and bottom three posts are stored in an online spreadsheet.

Using this method, we surfaced 1,022 overly shared or commented political posts and 272 coordinated links. In addition, we detected 66 coordinated political accounts (20 Facebook Pages and 46 public groups) and 554 generic coordinated accounts (406 Facebook Pages and 148 public groups) not listed in our initial accounts.

Findings

The examination of surfaced posts and identified accounts facilitated the discernment of various types of problematic information-sharing practices, orchestrated dissemination networks, and information operations that transpired during the 2022 Italian general election campaign. Some are new, while others are the evolution of practices already observed

(Giglietto et al., 2022). This section specifically concentrates on three networks that captured our attention as posts issued by accounts within these networks repeatedly surfaced. Although all exhibited some degree of coordination, these cases varied in terms of the types of actors involved (two were primarily composed of groups, while one consisted of Facebook pages) and the nature of their behaviors: one group aimed to sway the public electoral discourse in support of a specific political faction, another occasionally circulated misleading information for financial motives, and the third endeavored to promote religious conversion with spammy and inauthentic methods. Each case is analyzed using the A-B-C framework (François, 2019). Such a framework highlights the interplay between manipulative Actors, deceptive Behavior, and harmful Content. More precisely, the framework was tailored to our cases as follows: concerning the Actors, we examined the overall network's size and reach; with respect to Content, we concentrated on the types of content posted and their origins; ultimately, we investigated deceptive Behavior, assessing the dependability and bias of the disseminated sources and/or the potential distribution of problematic content.

Hyperpartisan Contents Circulated by M5S Groups

The first case concerns an unofficial group of accounts that supports the Italian populist party Five Star Movement–Movimento 5 Stelle – (M5S). Previous research (Giglietto et al., 2020a) has shown that M5S has historically relied on a large and active network of supporters who simultaneously post similar content across multiple public Facebook groups. This trend continued during the 2022 general election campaign, which revealed coordinated behaviors aimed at disseminating highly partisan content.

Size and Reach. Our pre-electoral CooRnet iteration featured a large cluster of coordinated entities (52) unofficially linked to M5S. As a result of the tracking activity performed by the workflow in the lead-up to election day, we detected 38 additional M5S-related entities. We have identified 90 entities (89 Facebook groups and 1 Facebook page) ascribed to the M5S network for a total potential reach of 1,547,159 users. Although these entities do not represent official party channels, their association with the M5S is evident from their names or descriptions, as they explicitly reference, for instance, the party's name or the leader, Giuseppe Conte. The top five pages by engagement in this cluster are "CONTE E CUORE IN MOVIMENTO" (Conte And Heart In Motion); "Raccolta firme per riportare al governo Giuseppe Conte" (A petition to bring Giuseppe Conte back to the government); "TELE TV 5 STELLE" (5 Star Tele Tv); "#Giuseppe Conte MoVimento 2050" (#Giuseppe Conte Movement 2050); "Amore per Conte ❤ Al tuo fianco Presidente !!" (Love for Conte ❤ At your side President !!).

Types of Content and Sources. Entities in this network were particularly prolific in terms of content posted: during the two months preceding the election day (July 21–September 25, 2022) they published 534,353 posts (81,410 or 20% with links), on the election day 6.2 posts/min on average with peaks of 50+ posts/min.

Using the full dataset of posts published by this updated list of accounts, we studied the sources circulated. Following a study on the "QAnon" conspiracy theory on Facebook (Kim & Kim, 2021), we examined the domains of the URLs included in these posts and classified the source into two main categories: Facebook Internal/Native and External. The first category (Facebook Internal/Native) refers to Facebook-native materials (photos, videos), accounts (pages, groups, and individual profiles), or services, while the second (External) refers to URLs that are not Facebook internal sources.[2]

Taken together, the large prevalence of posts without links (80%) and the majority of Facebook Internal Sources among posts with links (Table 1) suggest that External Sources play a minor role in shaping the information environment these users are exposed to.

Reliability and Hyperpartisanship. Taking into consideration the reliability and partisanship of sources, it is notable that the large majority of Facebook Internal Sources are native content (video, photos, etc.) created by personal profiles or hyperpartisan Pages and reshared in these groups (often by the same users who created the original post as a way to amplify the reach of their own creations). It is hard to assess the overall reliability of this content, and it is impossible to fully account for its recurrent pattern of circulation due to the limitation in the data provided by Meta. Despite these limitations and following Hindman et al.'s (2022) theory of "superusers," a limited number of very active members are responsible for circulating the same content to multiple groups in a very short period of time. The content published by these superusers is often shared on these groups by other personal profiles. When personal profiles do not create native Facebook content, it comes from a set of clearly partisan Pages which prominently feature the official Page of Giuseppe Conte and the party's official pages. The fact that the community and partisan Pages produce most of the recirculated content suggests that members tend to be eminently exposed to like-minded content.

On the contrary, posts with links to external sources can expose the members to alternative viewpoints, breaking the echo chamber's seals. We thus assessed the partisanship and reliability of the External Sources circulated by these accounts. Going into the entities classified as External Sources, about half of the URLs can be traced back to News Sources, that is, sources of information regularly registered and reliable (such as ilfattoquotidiano.it, lanotiziagiornale.it, and fanpage.it). The second largest category is Other Social Media Sources (e.g., YouTube.com, Instagram.com, etc.).

Table 1. Summary of the Classified Sources in Posts With Links.

Type of source	Freq	%
Facebook internal/Native sources	52,885	52.0
External sources	48,769	48.0
Total	101,654	100.0

Less than 10% of external sources are political/institutional, officially or unofficially connected to M5S. The categories Content Provider and Sexually Explicit Sources have low percentages, and everything else untraceable was classified as Other.

Using the NewsGuard rating (NewsGuard, 2020), we found that just 2% of the links come from known unreliable sources. However, due to the limited coverage of NewsGuard, most of the domains (76%) were not rated, hindering a proper third-party-supported reliability assessment of these sources. That said, it is worth noting that the large majority of news articles circulated belong to outlets considered close to the M5S (e.g., Il Fatto Quotidiano, NotiziaGiornale) or hyperpartisan blogs that aggregate links and/or reposts the full text of cherry-picked pro-M5S news stories.

Problematic Content. Aside from partisanship, monitoring the posts generated by these accounts during the campaign revealed at least two instances of problematic content that were reported, deemed false, and debunked by Italian fact-checkers. This problematic content involved the dissemination of fabricated polls, which allocated higher percentages of voters to the M5S than were found in the polls conducted. False polls are a manipulation technique frequently employed in electoral campaigns, as they aim to reinforce the bandwagon effect and enhance mobilization around a candidate (Natale, 2009). This can have even more disruptive consequences in highly partisan information environments like the one described above.

Premeditated Incidental Exposure to Misleading Political Content

The second investigated case concerns a network of Facebook Pages that coordinately posted clickbait image macros with links. The network was detected by investigating sources of posts that repeatedly surfaced for their performance and skewed comments/share ratio toward comments.

Size and Reach. The investigation traced a network of 46 Pages that published 58,035 posts from July 21 to September 25, 2022. Overall, the Pages belonging to this network claim to have various purposes, ranging from sharing comedy and aphorisms to religious pages that post prayers or image macros with biblical characters. The two accounts that achieve the highest levels of engagement within the network are two Pages that can be characterized as religious: "La Preghiera di Oggi" (Today's Prayer) and "Santa Rita da Cascia Avocada dei casi impossibili prega per noi" (Saint Rita of Cascia Avocada of Impossible Cases Pray for Us). The potential reach of these two pages was about 768,000 users since, at the time of analysis, they had 327,027 and 441,647 followers, respectively. They exchange intentions and primarily share image macros featuring biblical characters, prompting users to respond with prayers in the posts' comments.

Types of Content and Sources. Despite their claimed purposes, these Pages frequently contain news about current social and political issues in addition to their native entries. A considerable part of the published content is a mix of politics and entertainment, cross-shared by several Pages from the network. In this way, Facebook users looking for jokes, aphorisms, or prayers are also "incidentally" exposed to hard news.

The aforementioned religious Pages "La Preghiera di Oggi" (Today's Prayer) and "Santa Rita da Cascia Avocada dei casi impossibili prega per noi" (Saint Rita of Cascia Avocada of impossible cases pray for us), respectively, published 1,295 posts for total interactions (sum of comments, shares, and reactions including likes) of 3,799,049 and 2,366 posts (1,494,138 total interactions) between 21 July and 25 September 2022. Based on the name, description, and other self-presentation cues, both Pages claim to have a religious purpose. However, two-thirds of the posts are, in fact, not religious.[3] Despite their prevalence, non-religious posts tend to get fewer interactions than religious content (see Table 2).

Based on these data, we can ascertain that the vast majority of interactions are garnered by religious posts that are on-topic. However, this publishing behavior raises concerns because it potentially exposes a significant number of users—as indicated by Size and Reach—to clickbait image macros featuring links to news articles focused on political, occasionally controversial, issues alongside religious posts. In other words, while the low level of interactions highlights that most users subscribing to these Pages tend to be uninterested in non-religious posts, we can infer that many of them are indeed exposed to off-topic clickbait posts. Moreover, users who frequently interact with religious posts are more likely to be exposed to other posts as a side effect of the recommendation/distribution algorithm.

Reliability and Hyperpartisanship. The posts on these Pages frequently employ clickbait headlines that inadequately represent the news, directing the exposed users toward misleading or sensationalist interpretations.

Figure 2 illustrates the kind of news that the aforementioned pages coordinately publish and how the post appears on the user's feed.

Despite the extensive discussion on this news, they are frequently portrayed with the urgency typical of breaking news. All these posts share the same visual layout, composed

Table 2. Type of Posts and Average of Total Interactions (July 21–September 25, 2022).

		Religious posts	Non-religious posts
La preghiera di Oggi	Frequencies	394	901
	Total interaction	3,467,412	331,637
	Average of interaction	8,800.5	368.1
Santa Rita da Cascia Avocada dei casi impossibili prega per noi	Frequencies	607	1,759
	Total interaction	1,337,861	156,277
	Average of interaction	2,204.1	88.8

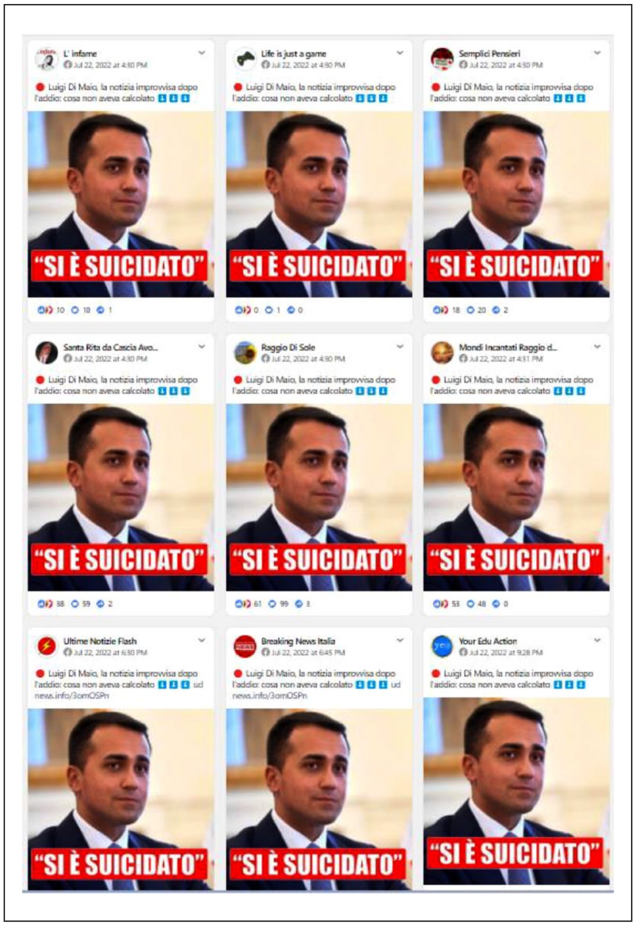

L' infame
Jul 22, 2022 at 4:30 PM
🔴 Luigi Di Maio, la notizia improvvisa dopo l'addio: cosa non aveva calcolato ⬇️⬇️⬇️

10 10 1

Life is just a game
Jul 22, 2022 at 4:30 PM
🔴 Luigi Di Maio, la notizia improvvisa dopo l'addio: cosa non aveva calcolato ⬇️⬇️⬇️

0 1 0

Semplici Pensieri
Jul 22, 2022 at 4:30 PM
🔴 Luigi Di Maio, la notizia improvvisa dopo l'addio: cosa non aveva calcolato ⬇️⬇️⬇️

18 20 2

Santa Rita da Cascia Avo...
Jul 22, 2022 at 4:30 PM
🔴 Luigi Di Maio, la notizia improvvisa dopo l'addio: cosa non aveva calcolato ⬇️⬇️⬇️

38 59 2

Raggio Di Sole
Jul 22, 2022 at 4:30 PM
🔴 Luigi Di Maio, la notizia improvvisa dopo l'addio: cosa non aveva calcolato ⬇️⬇️⬇️

61 99 3

Mondi Incantati Raggio d...
Jul 22, 2022 at 4:31 PM
🔴 Luigi Di Maio, la notizia improvvisa dopo l'addio: cosa non aveva calcolato ⬇️⬇️⬇️

53 48 0

Ultime Notizie Flash
Jul 22, 2022 at 6:30 PM
🔴 Luigi Di Maio, la notizia improvvisa dopo l'addio: cosa non aveva calcolato ⬇️⬇️⬇️ ud news.info/3omOSPn

Breaking News Italia
Jul 22, 2022 at 6:45 PM
🔴 Luigi Di Maio, la notizia improvvisa dopo l'addio: cosa non aveva calcolato ⬇️⬇️⬇️ ud news.info/3omOSPn

Your Edu Action
Jul 22, 2022 at 9:28 PM
🔴 Luigi Di Maio, la notizia improvvisa dopo l'addio: cosa non aveva calcolato ⬇️⬇️⬇️

Figure 2. Example of political news content by pages from the network.

of pictures in the background and a super-impressed clickbait headline poorly related to the attached news story. The link to the article itself, most of the time shortened, is either placed in the text that describes the photo-type post (message) or, occasionally, in the first comment of the post. In the example in Figure 2, the headline ("He killed himself") misleadingly implies the "suicide" of Luigi Di Maio, but it metaphorically refers to the Former External Relation Minister's political activity as "political suicide." The post may be intentionally misleading, as the metaphor is problematic for a less attentive reader.

Fact-checking this type of post is also rather complicated, as the term suicide, although misleadingly employed in this case, is sometimes used by Italian journalists to label political operations that turned out to self-inflict damages to the politician. Despite its misleading aim, the link points to Tvzap.it, a legally registered news magazine dedicated to celebrity gossip, a tactic that contributes to sowing further confusion in users encountering this news.

Problematic Content. The incidental exposure to news can have significant (and controversial) implications on how some people form their opinions. Some individuals may be more susceptible than others to the influence of news found randomly, depending on various factors such as age, level of education, cultural background, and political orientation. According to Fletcher and colleagues (Fletcher & Nielsen, 2018), people who do not actively seek out news are more affected by unintended news exposure as they have fewer resources at their disposal to form opinions or verify the accuracy of information. The incidental exposure to unreliable or false news may influence the perception of exposed subjects' reality, fueling phenomena such as intolerance, polarization, and conspiracy theories.

Religious Proselytism via Facebook Messenger Bots

The third and final case study, pursued based on surfaced posts, also underscores potential applications of the workflow concerning religious content. Contrary to the previous case, where religious Pages were utilized to disseminate problematic content, here we observe a reverse trend in which large public groups unrelated to religion serve as part of a complex operation to promote religious proselytism.

Size and Reach. While monitoring the campaign, the real-time analysis repeatedly surfaced well-performing posts with direct links to Facebook Messenger and WhatsApp chats posted by individual users in large Facebook public groups. A manual inspection of these posts and groups led to an ongoing ethnographic investigation on using Messenger bots to lure users and perform online proselytism activity. We traced this operation to the Church of Almighty God, also known as Eastern Lightning, a controversial religious movement established in

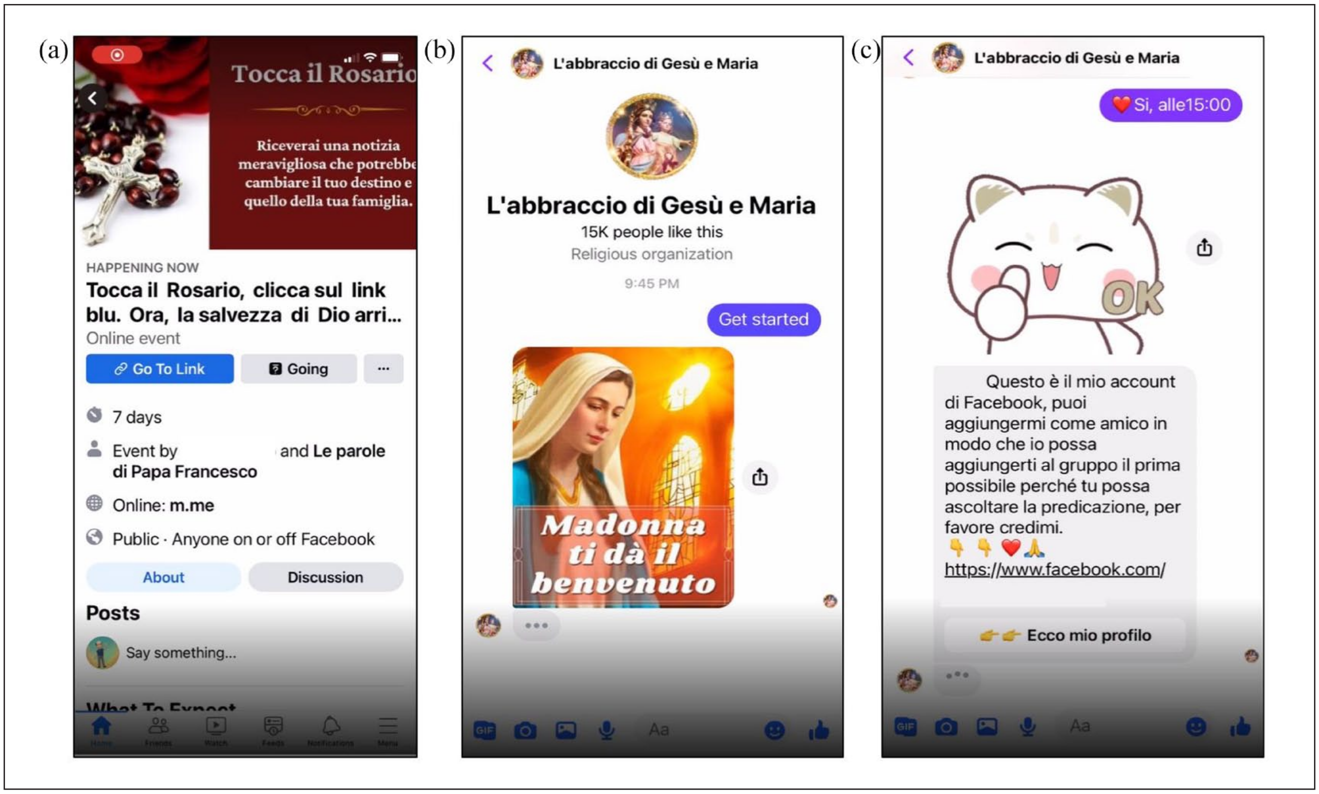

Riceverai una notizia meravigliosa che potrebbe cambiare il tuo destino e quello della tua famiglia.

Tocca il Rosario, clicca sul link blu. Ora, la salvezza di Dio arri...

Event by and Le parole di Papa Francesco

L'abbraccio di Gesù e Maria

❤️Sì, alle15:00

Questo è il mio account di Facebook, puoi aggiungermi come amico in modo che io possa aggiungerti al gruppo il prima possibile perché tu possa ascoltare la predicazione, per favore credimi.
👇👇❤️🙏
https://www.facebook.com/

👉👉 Ecco mio profilo

Figure 3. Screenshots exemplify the process initiated when the user clicks on the "Go to link" button advertised by the event. Personal profiles have been blurred.

China in 1991. To get a better picture of the Facebook assets involved in the operation, we collected all the posts created by all the monitored accounts posted between July 21 and September 25, 2022, and containing a Messenger link ("m.me") or a link to two domains affiliated to the operation "kingdomsalvation.org" and "vangelodigiorno.vangelodioggi.org." A CooRnet iteration on these links (N=13,696) detected a network of 1,390 public groups organized in three components and seven clusters grouped by language (two Italian, French, Spanish, English UK, English US, and Ukrainian). For this case study, we focused on the Italian groups (N=61 with 1,724,907 total members) and Pages (N=13 with 294,625 total subscribers) identified by analyzing the Messenger links.

Types of Content and Sources. The large, poorly moderated, and sometimes non-religious groups are used by the campaign to spread posts and events that bait users into their Messenger bot proselytism. The process starts with the user clicking the Messenger link (see Figure 3a). At this stage, the conversation is fully automated (Figure 3b). The purpose is to invite the user to join periodic catechism events in a semi-automated Messenger group chat. To get the invite, the user must first add a personal profile as a Facebook friend (see Figure 3c).

From a content point of view, the posts circulated by this network are generally photo-type posts containing a call-to-action caption that invites users to write "Amen" in comments, such as in Figure 4. Using this technique, these pages can reach tens of thousands of comments. The post shown in Figure 4 was published by the page "Vangelo di giorno" ("Gospel at day") with about 28,000 subscribers but reached over 90,000 comments.

Reliability and Hyperpartisanship. On average, the 13 Pages analyzed have 72.6 administrators each. This number is in itself remarkable. For example, the Page named *L'abbraccio di Gesù e Maria*[4] (The embrace of Jesus and Mary) has 122 admins, all Italians. Another Page, *Vangelo di giorno*[5] (The Daily Gospel) has 93 admins (91 Italian and 1, respectively, from Greece and South Korea). The Page was originally created on 15 March 2020 (during the COVID-19 first Italian wave) under the title *Andrà tutto bene* (Everything is gonna be fine), a popular slogan that emerged while the country was undergoing its strictest lockdown.

The high number of administrators and the misleading name changes is not the sole peculiarity of these entities' network. Another unusual characteristic is a high number of groups linked to the Page. An emblematic example is *L'amore sulla croce*[6] (Love on the Cross), which has 89 administrators and is linked to 79 Facebook groups.

Problematic Content. Overall, this network does not appear to be motivated by political or economic factors. Instead, the

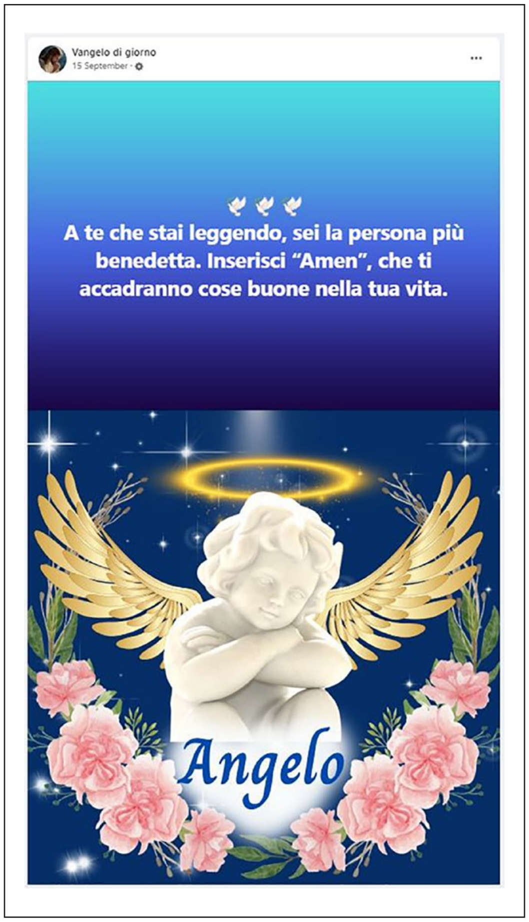

Vangelo di giorno
15 September ·

A te che stai leggendo, sei la persona più benedetta. Inserisci "Amen", che ti accadranno cose buone nella tua vita.

Figure 4. An example of these religious entities' posts. The caption is a call-to-action saying "To you who are reading, you are the most blessed person. Enter Amen, and good things will happen in your life."

primary goal of this large-scale operation is religious proselytism pursued through spammy and inauthentic strategies. Content shared by this network within large groups, often in the form of posts or events with links to their bots, never explicitly mentions the religious movement or refers to the aspects of their credo that differentiate them from mainstream Christianity, particularly the Catholic Church in Italy. Consequently, users drawn into their proselytism funnel are unaware that they are engaging with a specific religious movement and do not learn about the movement's true beliefs until they are already involved. At the time of writing, 8 of the 13 Pages investigated are unavailable.

Discussion, Implications, and Future Work

Information operations intent on manipulating public opinion by capitalizing on the vulnerabilities of social media platforms and the current media ecosystems are a persistent presence. While their peaks of activity spike during breaking news, crises, and elections, adversarial actors tend to be always active and ready to jump on whatever topic grabs the attention of a large enough audience. This article introduces a workflow to take advantage of these activity's peaks while enabling recurrent monitoring of problematic actors in their day-to-day audience-building and off-peak operations.

The workflow was implemented to follow the activities of known coordinated social media accounts during the 2022 Italian election. The content and actors that surfaced in the lead-up to election day prompted three investigations that shed light on the dynamic of attention-seeking performed by ideologically, economically, and religiously motivated groups.

Populists and hyperpartisan news outlets leverage existing social media assets to seek an audience for their content. Religious posts are used as an attention grabber to serve misleading political stories to an unaware audience. At the same time, inauthentic religious proselytism turns to large and poorly moderated Facebook groups to spread its message with spam techniques. In all three cases, the social media assets (Facebook Pages, groups, and personal profiles) belonging to each operation employed CLSB to amplify the reach of their content. Despite this common denominator, each case exhibits different problematic aspects and varies in size, reach, sources, and content type. Collectively, they help evaluate the workflow by demonstrating some outcomes of its implementation.

Despite these results, the implementation deployed is limited to two social media sites (Facebook and Instagram) and relies on a social media analytic infrastructure (CrowdTangle) that will reportedly be discontinued soon (Lawler, 2022). However, the workflow is designed to track known actors across different social media platforms by relying on the official researchers' APIs currently under development at Meta and already available for TikTok, Twitter, and YouTube. The detection logic can be easily adapted to any social media platform as long as their APIs support real-time post search. Post-search is the core building block of the architecture as it enables simple or more complex forms of near-duplicate detection (Papadopoulou et al., 2022) (e.g., posts with the same/similar link, message, image, audio, or video).[7] Such a form of data sharing is in accordance with Article 40 of the EU's Digital Services Act (DSA), which stipulates data access and scrutiny regulations for very large online platforms and search engines.

A limitation of this study is that it exclusively focuses on Italy's 2022 election campaign period. As highlighted, the value of this workflow lies in its capacity to facilitate monitoring beyond periods of heightened activity among coordinated actors. This is accomplished through continuous analysis over time, enabling the evaluation of differences between election and non-election periods. Consequently, the monitoring of out-of-peak activity periods is an ongoing endeavor and will be addressed in future research.

The implications of such an approach reverberate from the supplier to the consumer side of the equation. Constant monitoring of up-to-date lists of known actors may surface potentially problematic content in quasi-real time and feed a rapid-alert system for fact-checkers. An always-on strategy makes it easier to comprehend the tactics used by influence operations and reveals a broader, more recent, and comprehensive set of assets involved. Underestimating the assets employed could result in incorrectly estimating the campaign's reach and misclassifying the exposed people. A more current and thorough list of the assets involved can—instead—result in more accurate estimates of the prevalence of the campaign's content and, consequently, better identification of those who were exposed to assess the impact of that exposure on their behavior and opinions.

Finally, it is also worth noting that the workflow automatically highlights cases warranting further investigation, but neither the content nor the actors should be automatically deemed problematic or harmful on this basis. The alert system assists investigative journalists, researchers, fact-checkers, and debunkers in better prioritizing their work and allocating their limited resources to potentially problematic content gaining traction on mainstream social media. These expert investigations remain crucial for assessing the harmfulness of information operations.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This work was partially supported by the vera.ai project, which is funded by the European Union through the Horizon Europe program (Grant Agreement ID 101070093).

ORCID iDs

Fabio Giglietto https://orcid.org/0000-0001-8019-1035
Roberto Mincigrucci https://orcid.org/0000-0002-5754-387X
Anna Stanziano https://orcid.org/0000-0001-5525-9333

Notes

1. A post is considered political if its text fields matches at least one of these keywords: Fratelli d'Italia, FdI, Meloni, Partito Democratico, PD, Letta, Lega, Salvini, Movimento 5 Stelle, M5S, Conte, Forza Italia, FI, Berlusconi, Tajani, Azione, +Europa, Calenda, Italia Viva, Renzi, Italexit, Paragone, Alleanza Verdi Sinistra, Europa Verde, Verdi, Bonelli, Evi, Art.1-MDP, MPD, Speranza, Sinistra Italiana, Fratoianni, governo, parlamento, Draghi, Mattarella, elezioni, Impegno civico, Tabacci, Di Maio, Grillo. To minimize false positive political posts, we also double-check that some names of politicians that are also common Italian words (e.g., the surname of Giorgia Meloni is also the plural of a fruit name) are spelled with the first uppercase letter.
2. We identified six subgroups of external sources: other social media sources (domains of social media services), news sources (domains of registered news media), political/institutional sources (domains attributable to political parties or Italian institutions), content provider sources (domains attributable to content provider like), sexually explicit sources (domains that link to sexually explicit content), and other.
3. All posts that contained at least one of the following keywords in the post text or image were considered religious: amen, ave, benedica, bibbia, catechismi, catechismo, cattolica, cattolici, cattolico, cristo, crocefissi, crocefisso, crocifissi, crocifisso, dio, gesu, gesù, madonna, maria, medjugorje, pio, religione, religioni, religiosi, religioso, rosario, santo, vangeli, vangelo, vergine, preghiera.
4. https://www.facebook.com/Labbraccio.di.GesU.e.Maria/.
5. https://www.facebook.com/Vangelo.di.giorno.2019/.
6. https://www.facebook.com/sullacroce/.
7. Anonymize for peer review.

References

Allcott, H., & Gentzkow, M. (2017). Social media and fake news in the 2016 election. *The Journal of Economic Perspectives: A Journal of the American Economic Association*, *31*(2), 211–236. https://doi.org/10.1257/jep.31.2.211

Allcott, H., Gentzkow, M., & Yu, C. (2019). Trends in the diffusion of misinformation on social media. *Research & Politics*, *6*(2), 1–8. https://doi.org/10.1177/2053168019848554

Bastos, M. T., & Mercea, D. (2019). The Brexit Botnet and user-generated hyperpartisan news. *Social Science Computer Review*, *37*(1), 38–54. https://doi.org/10.1177/0894439317734157

Benkler, Y. (2019). *Cautionary notes on disinformation and the origins of distrust*. Social Science Research Council. https://doi.org/10.35650/md.2004.d.2019

Cinelli, M., Cresci, S., Quattrociocchi, W., Tesconi, M., & Zola, P. (2022). Coordinated inauthentic behavior and information spreading on Twitter. *Decision Support Systems*, *160*, 1–12. https://doi.org/10.1016/j.dss.2022.113819

Cresci, S., Di Pietro, R., Petrocchi, M., Spognardi, A., & Tesconi, M. (2016). DNA-inspired online behavioral modeling and its application to spambot detection. *IEEE Intelligent Systems*, *31*(5), 58–64. https://doi.org/10.1109/MIS.2016.29

Crowd Tangle Team. (2016). *How do you calculate overperforming scores?* https://help.crowdtangle.com/en/articles/2013937-how-do-you-calculate-overperforming-scores

Daniels, J. (2009). Cloaked websites: Propaganda, cyber-racism and epistemology in the digital era. *New Media & Society*, *11*(5), 659–683. https://doi.org/10.1177/1461444809105345

Davis, C. A., Varol, O., Ferrara, E., Flammini, A., & Menczer, F. (2016). BotOrNot: A system to evaluate social bots. In *Proceedings of the 25th International Conference Companion on World Wide Web* (pp. 273–274). https://doi.org/10.1145/2872518.2889302

Dawson, A., & Innes, M. (2019). How Russia's internet research agency built its disinformation campaign. *The Political Quarterly*, *90*(2), 245–256. https://doi.org/10.1111/1467-923x.12690

Donovan, J., Dreyfuss, E., Fagan, K., Faris, R., Friedberg, B., Holmes, C., Lim, G., Lytvynenko, J., Nilsen, J., O'Neil, M., & Salam, J. (2020). *The media manipulation definitions*. The Media Manipulation Casebook. https://mediamanipulation.org/definitions

Donovan, J., Dreyfuss, E., & Friedberg, B. (2022). *Meme wars: The untold story of the online battles upending democracy in America*. Bloomsbury Publishing.

Farkas, J., Schou, J., & Neumayer, C. (2018). Cloaked Facebook pages: Exploring fake Islamist propaganda in social media. *New Media & Society*, *20*(5), 1850–1867. https://doi.org/10.1177/1461444817707759

Fletcher, R., & Nielsen, R. K. (2018). Are people incidentally exposed to news on social media? A comparative analysis. *New Media & Society*, *20*(7), 2450–2468. https://doi.org/10.1177/1461444817724170

Forrester, B., Bacovcin, A., Devereaux, Z., & Bedoya, S. (2019). *Propaganda filters: Tracking malign foreign interventions on social media*. Nato. https://www.sto.nato.int/publications/STO%20Meeting%20Proceedings/STO-MP-IST-178/MP-IST-178-09.pdf

François, C. (2019). *Actors, behaviors, content: A disinformation ABC highlighting three vectors of viral deception to guide industry & regulatory responses (one)*. Transatlantic High Level Working Group on Content Moderation Online and Freedom of Expression. https://www.ivir.nl/publicaties/download/ABC_Framework_2019_Sept_2019.pdf

Fraser, L. (2020). *What data is CrowdTangle tracking?* CrowdTangle Help. https://help.crowdtangle.com/en/articles/1140930-what-is-crowdtangle-tracking

Freelon, D., McIlwain, C., & Clark, M. (2018). Quantifying the power and consequences of social media protest. *New Media & Society*, *20*(3), 990–1011. https://doi.org/10.1177/1461444816676646

Freelon, D., & Wells, C. (2020). Disinformation as political communication. *Political Communication*, *37*(2), 145–156. https://doi.org/10.1080/10584609.2020.1723755

Giglietto, F., Farci, M., Marino, G., Mottola, S., Radicioni, T., & Terenzi, M. (2022). Mapping nefarious social media actors to speed-up covid-19 fact-checking. *SocArXiv*. https://doi.org/10.31235/osf.io/6umqs

Giglietto, F., Righetti, N., & Marino, G. (2019). Understanding coordinated and inauthentic link sharing behavior on Facebook in the run-up to 2018 general election and 2019 European election in Italy. *Socarxiv*. https://doi.org/10.31235/osf.io/3jteh

Giglietto, F., Righetti, N., & Rossi, L. (2020). *CooRnet. Detect coordinated link sharing behavior on social media* (Version 1.0.0) [#rstat]. Github. https://github.com/fabiogiglietto/CooRnet

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

Giglietto, F., Righetti, N., Rossi, L., & Marino, G. (2020b). Coordinated link sharing behavior as a signal to surface sources of problematic information on Facebook. In *International Conference on Social Media and Society* (pp. 85–91). https://doi.org/10.1145/3400806.3400817

Giglietto, F., Terenzi, M., Marino, G., Righetti, N., & Rossi, L. (2020). *Adapting to mitigation efforts: Evolving strategies of coordinated link sharing on Facebook*. SSRN. https://papers.ssrn.com/abstract=3775469

Giglietto, F., Valeriani, A., Righetti, N., & Marino, G. (2019). Diverging patterns of interaction around news on social media: Insularity and partisanship during the 2018 Italian election campaign. *Information, Communication and Society*, *22*(11), 1610–1629. https://doi.org/10.1080/1369118X.2019.1629692

Gleicher, N. (2018, December 6). Coordinated inauthentic behavior explained [Interview]. https://about.fb.com/news/2018/12/inside-feed-coordinated-inauthentic-behavior/

Golebiewski, M., & Boyd, D. (2018). *Data voids: Where missing data can easily be exploited* (*Vol. 29*). Data & Society. https://datasociety.net/library/data-voids-where-missing-data-can-easily-be-exploited/

Graham, T., & QUT Digital Observatory. (2020). *Coordination network toolkit*. Queensland University of Technology. https://doi.org/10.25912/RDF_1632782596538

Gruzd, A., Mai, P., & Soares, F. B. (2022). How coordinated link sharing behavior and partisans' narrative framing fan the spread of COVID-19 misinformation and conspiracy theories. *Social Network Analysis and Mining*, *12*(1), 118. https://doi.org/10.1007/s13278-022-00948-y

Guess, A., Nagler, J., & Tucker, J. (2019). Less than you think: Prevalence and predictors of fake news dissemination on Facebook. *Science Advances*, *5*(1), 1–8. https://doi.org/10.1126/sciadv.aau4586

Guess, A., Nyhan, B., & Reifler, J. (2018). Selective exposure to misinformation: Evidence from the consumption of fake news during the 2016 US presidential campaign. *European Research Council*, *9*. https://pdfs.semanticscholar.org/a795/b451b3d38ca1d22a6075dbb0be4fc94b4000.pdf

Hindman, M., Lubin, N., & Davis, T. (2022, February 10). Facebook has a superuser-supremacy problem. *The Atlantic*. https://www.theatlantic.com/technology/archive/2022/02/facebook-hate-speech-misinformation-superusers/621617/

Houidi, Z. B., Scavo, G., Traverso, S., Teixeira, R., Mellia, M., & Ganguly, S. (2019). The news we like are not the news we visit: News categories popularity in usage data. *Proceedings of the International AAAI Conference on Web and Social Media*, *13*, 91–102. https://doi.org/10.1609/icwsm.v13i01.3212

Hristakieva, K., Cresci, S., Da San Martino, G., Conti, M., & Nakov, P. (2022). The spread of propaganda by coordinated communities on social media. In *Proceedings of the 14th ACM Web Science Conference 2022* (pp. 191–201). https://doi.org/10.1145/3501247.3531543

Integrity Institute. (2022). *Misinformation amplification analysis and tracking dashboard*. Elections Integrity Program. https://integrityinstitute.org/blog/misinformation-amplification-tracking-dashboard

Jack, C. (2017). *Lexicon of lies: Terms for problematic information* (Vol. 3). Data & Society. https://datasociety.net/library/lexicon-of-lies/

Jenkins, H. (2008). *Convergence culture: Where old and new media collide*. NYU Press.

Kim, S., & Kim, J. (2021). *Propagation of the QAnon conspiracy theory on Facebook*. OSF Preprint. https://doi.org/10.31219/osf.io/wku5b

Lawler, R. (2022, June 23). *Meta reportedly plans to shut down CrowdTangle, its tool that tracks popular social media posts*. https://www.theverge.com/2022/6/23/23180357/meta-crowdtangle-shut-down-facebook-misinformation-viral-news-tracker202262323180357

Magelinski, T., Ng, L., & Carley, K. (2022). A synchronized action framework for detection of coordination on social media. *Journal of Online Trust and Safety*, *1*(2), 1–24. https://doi.org/10.54501/jots.v1i2.30

Marino, G., & Giglietto, F. (2023). The power of Alternative Influence Networks (AIN) for spreading Covid-19 problematic information on Facebook during a year of pandemic. *Problemi Dell'informazione*, *2023*(1), 109–134. https://doi.org/10.1445/106772

Matatov, H., Naaman, M., & Amir, O. (2022). *Stop the [Image] steal: The role and dynamics of visual content in the 2020 U.S. Election Misinformation Campaign*. Arxiv. http://arxiv.org/abs/2209.02007

Mazza, M., Cola, G., & Tesconi, M. (2022). Ready-to-(ab)use: From fake account trafficking to coordinated inauthentic behavior on Twitter. *Online Social Networks and Media*, *31*, 1–10. https://doi.org/10.1016/j.osnem.2022.100224

Natale, P. (2009). *Attenti al sondaggio!* Laterza.

NewsGuard. (2020). *Rating process and criteria* [Internet Archive]. *NewsGuard*. https://www.newsguardtech.com/ratings/rating-process-criteria/

Nimmo, B., & Hutchins, E. (2023). *Phase-based tactical analysis of online operations*. Carnegie Endowment for International Peace. https://carnegieendowment.org/2023/03/16/phase-based-tactical-analysis-of-online-operations-pub-89275

Papadopoulou, O., Kartsounidou, E., & Papadopoulos, S. (2022). COVID-related misinformation migration to BitChute and Odysee. *Future Internet*, *14*(350), 1–22. https://doi.org/10.3390/fi14120350

Righetti, N., & Balluff, P. (2023). *CooRTweet: An R package to detect coordinated networks on Twitter* (Version 1.3) [Computer software]. https://cran.r-project.org/package=CooRTweet

Righetti, N., Rossi, L., & Marino, G. (2022). *At the onset of an infodemic: Geographic and disciplinary boundaries in researching problematic COVID-19 information* [Technical Reports]. Florida Marine Research Institute. https://doi.org/10.5210/fm.v27i7.12557

Santia, G. C., Mujib, M. I., & Williams, J. R. (2019). Detecting social bots on Facebook in an information veracity context. *Proceedings of the International AAAI Conference on Web and Social Media*, *13*, 463–472. https://wvvw.aaai.org/ojs/index.php/ICWSM/article/view/3244

Starbird, K., Arif, A., & Wilson, T. (2019). Disinformation as collaborative work: Surfacing the participatory nature of strategic information operations. *Proceedings of the ACM on Human-Computer Interaction*, *3*(127), 1–26. https://doi.org/10.1145/3359229

Tripodi, F. B. (2022). *The propagandists' playbook: How conservative elites manipulate search and threaten democracy*. Yale University Press.

Wardle, C., Greason, G., Kerwin, J., & Dias, N. (2018). Information disorder: The essential glossary. *First Draft*. https://firstdraft-news.org/wp-content/uploads/2018/07/infoDisorder_glossary.pdf

Wijffels, J. (2022). *cronR* (Version 0.6.2). https://github.com/bnosac/cronR

Wu, B., Liu, L., Yang, Y., Zheng, K., & Wang, X. (2020). Using improved conditional generative adversarial networks to detect social bots on Twitter. *IEEE Access*, *8*, 36664–36680. https://doi.org/10.1109/ACCESS.2020.2975630

Yang, Y., Davis, T., & Hindman, M. (2023). Visual misinformation on Facebook. *The Journal of Communication*, *73*, 316–328. https://doi.org/10.1093/joc/jqac051

Zuckerberg, M. (2021). A blueprint for content governance and enforcement. *Facebook*. https://www.facebook.com/notes/75-1449002072082/

Author Biographies

Fabio Giglietto (PhD "University of Urbino") is an Associate Professor of Internet Studies at the University of Urbino Carlo Bo. His key research focuses on information theory, societal impacts of media and digital technologies, and their interplay with social systems.

Giada Marino (PhD "University of Urbino") is a Postdoctoral Researcher at the University of Urbino Carlo Bo. Her research examines the relationship between information disorder and political polarization, with a specific emphasis on citizen engagement on social media platforms.

Roberto Mincigrucci (PhD "University of Perugia") is a Postdoctoral Researcher at the University of Urbino Carlo Bo. His research interests focus on mediated scandals, the relationship between journalism and corruption, and the analysis of the different forms of political communication.

Anna Stanziano (PhD "University of Perugia") is a Postdoctoral Researcher at the University of Urbino Carlo Bo. Her main research interests are political communication, journalistic coverage of corruption, the perception of corruption, and the relationship between media and religion.