POLITECNICO DI MILANO

Facoltà di Ingegneria

Scuola di Ingegneria Industriale e dell'Informazione

Dipartimento di Elettronica, Informazione e Bioingegneria

Corso di Laurea Magistrale in

Computer Science and Engineering

Towards Understanding Alternative and Mainstream News Dissemination on Social Media: The case of the 2018 Italian Election

Advisor:

PROF. STEFANO ZANERO

Co-Advisor:

PROF. GIANLUCA STRINGHINI

DOTT. MICHELE CARMINATI

Tesi di Laurea Magistrale:

SIMONE PATUELLI

Student Id n. 899290

Academic Year 2019-2020

Alla mia famiglia, per il sostegno, supporto, e carica che mi avete dato durante tutto questo mio percorso. Siete la mia fonte di energia.

ACKNOWLEDGMENTS

Mi sembra giusto dedicare questo spazio del mio elaborato alle persone che hanno contribuito, con il loro supporto, alla realizzazione dello stesso. Cercherò di essere breve quindi mi vogliano scusare tutte le persone che non citerò.

In primis, voglio ringraziare il mio relatore, Stefano Zanero. Quando mi sono presentato nel suo ufficio per ottenere più informazioni riguardo le possibilità di tesi, non mi sarei mai aspettato un'esperienza del genere. Grazie per il lavoro scelto, per tutto il supporto, e per avermi concessso di svolegerne parte in USA. Un secondo enorme ringraziamento va al professor Gianluca Stringhini. Per avermi ospitato nel proprio laboratorio a Boston e avermi fatto lavorare su un proprio progetto. Mi avete permesso di realizzare un sogno.
Ringrazio il professor Michele Carminati, per il sostegno, gli indispensabili consigli e aiuti che mi hai concesso durante tutto il lavoro. Senza di voi la mia tesi non sarebbe stata possibile.

Ringrazio infinitamente la mia famiglia che mi ha sempre sostenuto, appoggiando ogni mia decisione e ambizione. Avete creduto in me prima che iniziassi a crederci io stesso. Vi sarò eternamente grato. Un grazie speciale a Samantha, per avermi sopportato e supportato durante tutto il percoso di studi. Con te ho superato i momenti più difficili e senza non ce l'avrei mai fatta.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

ABSTRACT

In recent years, the patterns with which news is produced and consumed have changed. Nowadays, anyone can write news and find a channel to push it on the web. Furthermore, social media have became the first place to look for news to the point that, in 2017, two-thirds of Americans reported getting their news from social platforms. However, the ease with which misleading content can gain visibility on social networks has made them a fertile ground for the spreading of misinformation. Moreover, for the implementation of disinformation campaigns. The alleged Russian interference in the 2016 United States elections is just an example of what is happening in many countries. Several investigations in the Italian panorama have brought to light the presence of massive alternative media networks. They were composed of Facebook pages, news websites, and social accounts of politicians and they share mostly misleading content close to the 2018 Italian general elections.
In this thesis we perform a deeper analysis on the phenomenon of the dissemination of alternative and mainstream news on social media. In particular, we analyze the Italian scenario during the 2018 election period, focusing our attention on a set of Facebook suspicious accounts. The results obtained show that the 2018 Italian elections have seen the use of social campaigns intended to share large quantities of disinformation. Alternative and mainstream news have been shared with different temporal patterns. For example, alternative news have been highly posted close to important events aiming at influencing the public opinion. Finally, using the Hawkes Processes statistical model, we compute the influence that the studied accounts had on pushing URLs on Twitter. We find that those Facebook accounts were more influential at disseminating news URLs towards Twitter than a general Facebook dataset. This result is emphasized if we consider only mainstream news. Besides, the studied Facebook accounts tended less likely to be influenced by the URLs shared on Twitter.

SOMMARIO

Il rapido avanzamento di Internet nell'era moderna ha portato ad un cambiamento considerevole nel ramo dell'informazione. I media tradizionali come radio, televisione e giornali, sono stati sostituiti dalle rispettive versioni digitali. Inoltre, la rapida diffusione negli ultimi anni delle piattaforme social ha dato forma ad uno scenario completamente nuovo; l'informazione è direttamente prodotta e consultata attraverso i social networks. Uno studio, condotto da Pew Research Center, riporta che nel 2017 due terzi degli Americani hanno tratto le proprie notizie dai social media [36].

Il cambiamento nei metodi con cui l'informazione è distribuita ha modificato la natura stessa delle notizie. In passato, l'informazione veniva riportata solo da giornalisti professionisti del proprio settore. Oggigiorno invece, chiunque può pubblicare un articolo, senza preoccuparsi che abbiano un fondamento reale, e trovare un canale attraverso cui diffonderlo in rete. I social media sono strutturati in modo tale che un post, come può essere un articolo, abbia semplicemente bisogno di likes e condivisioni per diventare popolare. Si potrebbe, dunque, dedurre che quest'aspetto dei social media è la ragione per cui titoli accattivanti e di basso livello vincono l'attenzione del pubblico. Il timore che questo meccanismo possa facilitare il diffondersi di informazioni ingannevoli è stato confermato da un recente studio condotto da un team del Massachusetts Institute of Technology. Il leader del team ha, infatti, dichiarato che "il diffondersi di falsità supera quello verità perché l'essere umano è più propenso a ritwittare notizie false che non vere," [41].

La facilità con cui si può trovare misinformazione online, non rappresenta l'unico rischio riferito al consumo di notizie sulle piattaforme social. Infatti, datone l'enorme potere comunicativo, oggigiorno i social media rappresentano il principale strumento di comunicazione utilizzato dai politici di tutto il mondo. Secondo un altro studio condotto da Pew Research Center, il 44% degli Americani hanno ammesso di aver ottenuto le informazioni riguardanti le elezioni presidenziali del 2016 da piattaforme social [5]. Ecco quindi che la struttura aperta e connessa, sui cui i social media sono basati, ha permesso l'attuazione di campagne di disinformazione. Queste operazioni sono pianificate nel dettaglio e vengono svolte da utenti malevoli che hanno il solo scopo di manipolare l'opinione pubblica. Un esempio è rappresentato dalla campagna svolta dai Russi con lo scopo di interferire con le elezioni americane del 2016. Attraverso la creazione di milioni di accounts fasulli e promuovendo notizie tendenziose o inventate, hanno infatti sostenuto la candidatura di Donald Trump [37, 26].

Gli Stati Uniti, però, non sono stati gli unici paesi ad aver visto in prima persona l'attuarsi di campagne di disinformazione. Infatti, uno studio [3] da parte della Università di Oxforford ha evidenziato che quest'ultime, oggigiorno, rappresentano una minaccia mondiale.
In passato nel panorama italiano, grazie ad alcune indagini svolte da media e giornalisti, erano iniziate a circolare voci riguardo la distribuzione di informazioni fasulle. Il fine, si intendeva fosse quello di manipolare l'opinione pubblica riguardo eventi o partiti politici. Nel novembre del 2016, Buzzfeed pubblicò un articolo [27] con il quale accusava il Movimento 5 Stelle di aver creato una rete attorno al quale si snodava un'intensa attività di diffusione di notizie false e teorie cospirazioniste. La rete disinformativa comprendeva, oltre agli account social e il blog del partito stesso, un vasto assortimento di siti di notizie che si affermavano indipendenti. Ad un anno di distanza, lo stesso Buzzfeed pubblicò un nuovo rapporto [28] rivelando l'esistenza di un network contenente un vasto numero di notiziari online e pagine Facebook che diffondevano retorica nazionalistica, contenuti anti-migranti e disinformazione.
Ecco quindi che alla luce delle elezioni Italiane del 2018, non sorprende affatto la frase rilasciata in un intervista al New York Times da Matteo Renzi. L'ormai ex Primo Ministro, dichiarava: "Noi ci rivolgiamo ai social ma specialmente a Facebook per chiedere un aiuto per una campagna elettorale chiara e pulita" [16].

Alla luce della presenza di rumors sui contenuti condivisi durante le elezioni Italiane del 2018, la nostra ricerca propone un'analisi approfondita sul fenomeno della diffusione di notizie mainstream e alternative sui social media.
Per affrontare questo studio, focalizziamo la nostra attenzione sui post di 23 account Facebook. Questi account ci sono stati segnalati da un giornalista esperto nella verifica dei fatti a causa dell'alto contenuto misinformativo condiviso. Oltre ai post Facebook, ci siamo occupati di creare un secondo dataset basato su Twitter. Questo, è composto da tweet contenenti le stesse parole chiave e domini estratti dal dataset di Facebook. Con l'obiettivo di poter qualificare il contenuto condiviso, abbiamo classificato i siti di notizie che compaiono dei post raccolti in due categorie; alternativi e mainstream. Il processo di classificazione utilizza solo strumenti pubblicamente disponibili e si basa principalmente sul punteggio attribuito da NewsGuard[1]. Facciamo, inoltre, utilizzo di alcune liste contenenti siti web segnalati per la bassa qualità dei contenuti condivisi [2] [3].

1 https://www.newsguardtech.com/
2 https://www.bufale.net
3 https://www.butac.it

Successivamente, analizziamo la dinamiche temporali dei post raccolti. Per avere un idea sugli eventi importanti avvenuti durante il periodo elettorale, in questa parte dell'analisi prendiamo come punti di riferimento principalmente due date: Il giorno delle elezioni - 04/03/2018 - e il giorno in cui il governo Conte è entrato in carica - 01/06/2018. Inizialmente, studiamo la variazione del volume di post condivisi durante il periodo elettorale del 2018. Miriamo a vedere come il comportamento degli account si è evoluto nel tempo in relazione alla condivisione di notizie mainstream e alternative. Dopo, analizziamo come i due tipi di contenuti sono consumati su Facebook e Twitter. La velocità con cui le notizie sono ri-condivise e con cui appaiono da una piattaforma all'altra, sono alcuni dei parametri che prendiamo in considerazione.
Successivamente, passiamo all'analisi del contenuto testuale dei post. Qui, sfruttando tecniche di Sentiment Anlaysis, studiamo le dinamiche emotive estratte dai post raccolti. Approfondiamo anche lo studio dei sentimenti in relazione ai maggiori esponenti politici del tempo. Dopo, analizziamo quali hashtags sono soliti accompagnare le notizie mainstream e quali le notizie alternative.
Essendo l'obiettivo finale delle campagne disinformative quello di influenzare il maggior numero possibile di persone. E' ragionevole presumere che l'attività degli account Facebook miri a raggiungere altre piattaforme social. Utilizzando un approccio statistico chiamato Hawkes, abbiamo stimato l'influenza che l'attività degli account Facebook ha avuto su Twitter. L'idea è quella di considerare le due piattaforme social come collezioni di eventi. Quando un evento avviene su un social - la condivisione di una notizia - è causata una risposta ad onda che incrementa la probabilità di una ri-condivisione della stessa notizia sul social stesso, o sulle altre piattaforme social considerate. Per ogni notizia, siamo quindi in grado di calcolare se un social ha influenzato la sua condivisione o meno. Considerando tutti gli eventi raccolti, possiamo ottenere un punteggio che descriva quanto l'attività su Facebook ha influenzato quella su Twitter e viceversa.
Il nostro lavoro ha portato alle seguenti osservazioni.

1) Le elezioni italiane del 2018, come altri eventi politici prima nel mondo [11, 37], hanno visto l'uso di campagne sociali intese a condividere grandi quantità di mininformazione. Questo risultato è confermato dal fatto che l'attività della maggior parte degli account Facebook studiati è innescati dall'avvicinarsi delle elezioni. Infatti, diventano particolarmente attivi nel promuovere contenuti alternativi solo per quel periodo specifico. 2) Gli account Facebook studiati utilizzano spesso schemi di condivisione diversi per contenuti alternativi e mainstream. Ad esempio, mentre le notizie alternative sono fortemente condivise sia prima che dopo eventi importanti. Le notizie mainstream, invece, vengono condivise spesso solo dopo. Inoltre, sia

su Facebook che su Twitter, le notizie alternative sembrano essere più condivise a tardi orari. 3) I nostri studi sull'influenza mostrano che gli account di Facebook analizzati sono più efficienti nell'influenzare il dataset di Twitter, rispetto al comportamento di un dataset di Facebook generale. Inoltre, gli account Facebook tendono ad essere meno influenzati dall'attività degli utenti Twitter, comportandosi come un sistema chiuso in entrata. 4) L'influenza degli account Facebook verso Twitter è più marcata per la categoria di notizie mainstream. Infatti, questo tipo di notizie risulta essere più efficiente nel causare eventi su Twitter rispetto alle notizie alternative. Questo risultato è in contrapposizione con quello ottenuto per il dataset Facebook generico. Da quest'ultimo infatti, sembrerebbe che Facebook sia generalmente più efficiente a influenzare Twitter con notizie alternative.

1 INTRODUCTION

The rapid advance of the Internet in the modern era has led to a considerable shift in the information branch The attention was moved from traditional media outlets such as radio, television, and newspapers, to digital ones like news outlet websites and collectors. Moreover, the widespread diffusion of social platforms in the last years has led to a further turn, shaping a new scenario, where information is directly produced and consumed through these platforms. A study conducted by the Pew Research Center found out that in 2017 two-thirds (67%) of Americans reported getting their news from social media [36].

The shift of the media patterns through which information is being shared has also changed the nature of the news itself. While in the past, news was written by journalists experienced in their respective fields, nowadays anyone can immediately publish news, whether the information is reliable or not, and find a channel to push it on the web. Social media are structured so that a piece of information, such as an article, merely needs to be "liked" and shared multiple times to gain enough attention to become visible. One can argue that this social media aspect is the reason of why sensational and ridiculous headlines make it to the front, where it wins the competition for attention. The fear that this new mechanism is facilitating the dissemination of misleading information is confirmed in a recent study conducted by a team of the Massachusetts Institute of Technology on the journal of Science. They state “, the spread of falsity is outpacing the truth because human beings are more likely to retweet false than true news,” [41]. The emergence of such a wide, heterogeneous, mass of information sources, and the presence of unsubstantiated or untruthful rumors on social media, is contributing to the alarming phenomenon of misleading information [47].

Disinformation is generally used to refer to deliberate (often orchestrated) attempts to confuse or manipulate people through delivering dishonest information to them. *Misinformation* is used to refer to misleading information created or disseminated without manipulative or malicious intent. Both are problems for society, but disinformation is particularly dangerous because it is frequently organised, well resourced, and reinforced by automated technology [17].

Furthermore, the spread of misinformation is not the only risk related to the consumption of news on social media. Given the enormous communicative power carried out by social media platforms, suffice to say that in 2019 Facebook alone had 2 billion monthly active users [1], they have been largely utilized by the world's politicians. According to the Pew Research Center 44% of the Americans have admitted to getting their information regarding the 2016 presidential election from social media platforms [5]. Hence, the open and connected structure of social media allows the implementation of disinformative campaigns. This operations are planned in details and organized by malicious users who aim to manipulate public opinion. Through disinformation campaigns social media are being weaponized. An example of this behavior is represented by the campaigns carried out by Russian trolls in order to interfere with the 2016 American elections, boosting Donald Trump candidacy [37, 26]. Researches have proven evidences that political opinions have been manipulated by millions of fake social media accounts, leading ad hoc campaigns of disinformation, promoting messages that were slanted or even made up [26].

Disinformation campaigns are not targeting only US. A report from the Oxford University evidences the fact that disinformation campaigns are, nowadays, a global problem [3]. The report contains a list of 70 countries in which it is been proven the use of bots, fake social media accounts, and hired "trolls", to perform computational propaganda. Among this list, it is possible to find Italy too. It is reported for the use of the previous-cited tactics with the aim to distract or divert conversations away from important issues and drive division and polarization. In the Italian panorama, rumors related to the dissemination of disinformation to manipulate the public opinion had already come to light thanks to several investigations made by media and fact-checkers. On November 2016, *Buzzfeed* released an article [27] accusing the Five Star Movement (M5S) to have built a sprawling network of websites and social media accounts that were spreading fake news, conspiracy theories, and pro-Kremlin stories to millions of people. This network included not just the party's own blogs and social accounts, which have millions of followers, but also a collection of websites that describe themselves as "independent news" outlets. Those websites mainly shared M5S campaign lines, misinformation, and attacks on political rivals using sensational headlines as: "THE TRUTH THEY ARE TRYING TO HIDE FROM US". Exactly one year later, *Buzzfeed* released a new report [28]. This time exposing a massive network of Italian news websites and Facebook pages that spread nationalist rhetoric, anti-migrant content, and misinformation. This network represents one of the most popular alternative media operations in Italy. In fact, their content received more than 5 million shares on the social platform over a year. It comes with no surprise the statement "We ask the social networks, and especially Facebook,

1 https://urly.it/35da6

to help us have a clean electoral campaign," released by Matteo Renzi, ex Prime Minister, in an interview [16] for the New York Times few months before the new general elections of 2018.

In light of the presence of misleading content in the 2018 Italian elections, our work performs a deeper analysis on the phenomenon of the dissemination of alternative and mainstream news on social media. We focus our attention on the dynamics happening on Facebook and Twitter. For Facebook, we analyse the posts of 23 suspicious accounts. They were originally reported to us by a journalist experienced in fact checking. Then, we create a Twitter dataset made of tweets containing the same keywords and domains included in our Facebook dataset. We, therefore, want to understand if, from those data, it is possible to confirm the presence of disinformation campaigns during the general election of 2018. We aim to understand the differences of behaviour between the spreading patterns of alternative and mainstream news, how the studied accounts operate, and the influence exercised by them towards Twitter.
Our main findings indicate that:

1. The 2018 Italian elections have seen the usage of social campaigns intended to share large quantities of misinformation. This behavior is confirmed by the fact that most of the accounts in the Facebook dataset are triggered by the approaching of the elections and become very active in pushing alternative content just for that specified period.

2. Alternative and mainstream news are posted with important differences over time, showing unique patterns. For example, on Facebook, while alternative news are heavily pushed both before and after events, mainstream news are mainly posted after. Moreover, on both Facebook and Twitter, alternative news are spread more during late afternoon.

3. Our influence estimation experiments reveal that the Facebook accounts studied are more influential at disseminating news URLs towards Twitter than a general Facebook dataset. Besides that, those accounts tended less likely to be influenced by Twitter activity.

4. Looking at the specific type of content spread, we have seen that the Facebook accounts studied are more efficient at disseminating mainstream news rather than alternative one. This result is in opposition with the baseline, that show that generally Facebook is more efficient at pushing alternative content towards Twitter.

2 BACKGROUND AND MOTIVATION

In this section we provide a brief overview of the social media platforms we study. Then, we review previous work on the dissemination of news, misinformation spreading, and politically motivated campaign of disinformation on social media. These topics have become popular over the last couples of years. That is why there are many different studies and researches that face those problems from different point of views and using many approaches. We focus on works strongly related to our goal.

2.1 PLATFORMS BACKGROUND

Our work focus on Facebook and Twitter. We choose them mainly for three reasons: 1) On the Italian panorama, when it comes to social media, Facebook head the ranking, as it was used by 90.4% of online users in 2018 [1]. Twitter ranked fourth with 23.8% of online users [1]. But considering only text based social platforms, Twitter it is second only to Facebook. 2) Previous studies pointed out Facebook as the dominant platform for disinformation campaigns [3] 3) As far as we know, we are the first to conduct our analysis and consider the relationship between this two platforms.

FACEBOOK [2] is a social network where users can post comments, media and links to news or other content on the web, chat live, and watch videos. Facebook allows members who have common interests to find each other, interact in groups, and promote a public page built around a specific topic. It is also possible to publicize an event with the additional possibility to invite guests and track who plans to attend. Facebook is the largest social media service on the planet, with over 2 billion monthly active users[3] — over a quarter of the world's population. In 2019, the top 100 fake news stories on Facebook were viewed over 150 million times[4] .

TWITTER [5] is a micro-blogging directed social network on which users post and interact with messages known as "tweets". Users can post, like, and share (re-tweet) tweets made by the accounts

1 https://urly.it/35dj7
2 https://facebook.com
3 https://urly.it/35dj5
4 https://urly.it/35dj6
5 https://twitter.com

they follow or by public accounts. Some of its features include the hashtag (a keyword preceded by #), which makes it easier for users to find and weigh in on tweets around a theme. In January 2020 Twitter counts 330 million monthly active users[3].

2.2 RELATED WORK

A wide literature branch is devoted to understanding the information propagation dynamics on socials platforms, trying to determine the way in which news gets consumed and which characteristics make it viral. Bakshy et al. [2] analyse the information spread on Facebook related with user to user influence. Using a large-scale experiment that randomizes the exposure to signals, they prove that those users that are highly exposed to links are significantly more likely to spread information. They also show that weak ties, in users networks, are those responsible for the propagation of novel information. Lerman et al. [20] conduct a study on Digg and Twitter. They prove that the structure of the underlying networks highly influence the dynamics of information spreading on the platforms themselves. Zannetou et al. [44] put the foundation to the understanding of how mainstream and alternative news flows between Twitter, Reddit, and 4chan. Collecting millions of post from the three different social platforms, they studied events coming from 99 different news site, which they have previously classified as either mainstream or alternative. Using mainly cross platform analyses they are able to get a deeper understanding of the connections that link the different social media communities. Next, using the Hawkes processes [34], they modeled the influence between platforms in order to have an idea about the degree of influence on each other. It came out that Twitter is actually influenced by smaller fringe Web communities, for example the sub-reddit called The_Donald.

More recent studies have shifted their attention on the problem of disinformation spreading, focusing on the propagation dynamics in social networks. Zannetou et al. [46] study the behaviour and posted content of state-sponsored trolls, as well as their influence on different social platforms. They used a dataset, released by Twitter on 2018, of Russian and Iranian trolls accounts. They also gather data from Reddit using a list, released by Reddit itself, of Russian trolls accounts. Using word embedding [39], they found out that Russian trolls were pro-Trump while Iranian ones were against him. Analysing their behavior during the whole time period, they found out that trolls that used to post in a specific language tend to change it near real world events, as effort to be more influential over public opinion. Finally, analysing the influence through Hawkes processes [34], their study show that both types of trolls are particularly influential in pushing content to

Twitter. Shao et al. [35] created Hoaxy, a platform for the collection, detection, and analysis of online misinformation and its related fact-checking efforts. They also analyse a sample of tweets finding out that fact-checking content typically lags that of misinformation by 10-20 hours.

Moving to the Italian panorama, the project of Giglietto et al. [13] aims to create a mapping of the media coverage on political issues, produced by the Italian media, during the Italian elections of 2018. Their work measure the volume of interactions produced around news on Facebook and Twitter, estimate the political trend of the different sources, evaluate the polarization level of online audience and analyze in depth three cases of problematic information that had a significant impact on the campaigns. Using an approach called "Multi Party Media Partisanship Attention Score" [14], they are able to estimate a score towards a political party for each news source. Zollo et al. [47] investigate how information related to two very distinct narratives - scientific and conspiracy news - gets consumed on Facebook. The former category consider pages posting scientific knowledge and it includes institutions, organizations and scientific press. Pages posting contents neglected by mainstream media are selected for the second category. Their goal is to study users interactions with the Facebook pages belonging to such categories, between 2010 and 2014, in both the Italian and US context. Looking at the engagement of each user toward a specific type of content, they found out the existence of echo chambers, where users interact with like-minded people sharing their own system of beliefs. Measuring the response to the injection of deliberately false information, they discovered that users belonging to conspiracy pages are more likely to jump the credulity barrier, producing an higher number of likes and comments to troll posts.

As just mentioned, many previous studies focused on the field of misinformation and disinformation spreading on social platforms. However, to the best of our knowledge, a real scenario of disinformation diffusion in the Italian panorama has never been studied more in details.

2.3 RESEARCH GOALS

Unlike previous work, this research focus on a set of Italian Facebook accounts, reported for the unusual quantity of misinformative content shared on the platform. We analyse the patterns used by them in the dissemination of alternative and mainstream news. We, therefore, want to understand if, from those data, it is possible to respond to some important questions such as: Can we confirm the presence of disinformation campaigns, in the Italian panorama, during the general election of 2018? How do these accounts operate? Do alternative and mainstream news follow similar patterns? Are URLs matched with particular content inside posts? Can we quantify the influence that the accounts have on external platforms ?

3 APPROACH

In this section we provide a description of the approach we propose to answer the previous questions. We describe the process of data collection, the pipeline used for the classification of news domains in mainstream or alternative, and the main steps done in the analysis.

3.1 GENERAL OVERVIEW

In this research we study the news dissemination patterns followed by a set of Facebook suspicious pages, focusing in the period around the 2018 Italian elections. To do that, we extract the domains from the shared news and, using a semi-automatic process, we classify them in mainstream or alternative. To validate our classification process we do a brief cases study of some of the domains with the highest and lowest score. Next, we crawl the 1% of Twitter in order to have an additional set of "normal" users and behaviors to compare with. We analyse the temporal dynamics and the posts content mainly using data visualization approaches. At each step, we compare the results obtained for the two categories of news in the different platforms. Finally, using a statistical approach, we look at how the dissemination of news on Facebook influenced Twitter. Note that in our influence estimation experiments, we use an additional Facebook dataset that managed to fit the problem, as a baseline for a more generic Facebook behavior.

3.2 APPROACH DETAILS

This section contains the details of all the main steps done during the research.

3.2.1 *Collecting Facebook Data*

We create our main dataset collecting Facebook posts for around a year until October 10, 2018. The script focused on 23 pages which shared content strongly related to the Italian Election of 2018. For the part of estimating the influence on other social platforms, we also used an additional Facebook dataset, publicly available[1]. This datasets consists of observations of the Facebook engagement around Italian political news.

1 https://urly.it/360h2

From now on, we refer the datasets containing the posts of the 23 accounts by calling it just *Facebook*, instead, when we use the second Facebook dataset we call it *Facebook general*.

3.2.2 *Collecting Twitter Data*

We deploy an infrastructure which collects the 1% of publicly available tweets, every day, using the Twitter Streaming API[2]. From the stored tweets, we looked for those containing an URL, posted from January the 1st 2017 to October the 1st 2018, having:

- One of the keywords extracted from the posts on our *Facebook* dataset.
- One of the hashtags extracted from the posts on our *Facebook* dataset.
- An URL coming from one of the valid domains shared in the posts on our *Facebook* dataset or on the *Facebook general* dataset.

In this way we collected 115 497 tweets. Since tweets are collected at the time they are posted, we don't have information about the number of re-tweets or likes.

3.2.3 *Classification of news domains*

Next, we create a semi-automatic pipeline which aim to classify domains in either alternative or mainstream, using public tools. In particular, we exploit Virus Total[3], NewsGuard[4], Butac.it[5], and Bufale.net[6].

VIRUS TOTAL is an online service that aggregates many antivirus products and online scan engines to check for viruses on uploaded files. Virus Total can also be used as a searching engine, for domains or URLs, to retrieve the information it aggregates on the checked domain or URL. For us it is important that the report returned contains one or more categories that describe the domain's content. The categories returned in the report are gathered by Virus Total from different engines like Alexa, BitDefender, and TrendMicro.

NEWSGUARD: It is a website created in 2018 which keeps track of news sites and show for each of them a score. Trained journalists

2 https://urly.it/3610t
3 https://www.virustotal.com
4 https://www.newsguardtech.com/
5 https://www.butac.it
6 https://www.bufale.net

and experienced editors, employed by NewsGuard Technologies, review and rate news sites on their reliability and general trustworthiness. Giving a list of nine journalistic criteria, each of them worth a specific number of points, sites that score at least 60 out of 100 points display a green icon next to their name. The ones that score lower than 60, instead, get a red icon. Some of the NewsGuard criteria to determine a website's overall credibility score include:

- The frequency of publication of inaccurate information.
- The extent of sourcing and original reporting of information.
- The degree of demarcation between news and opinion journalism.
- The accuracy of headlines, including the use of click bait headlines.
- The degree of disclosure of the website's ownership, as well as the political positions of the owners.

Even if NewsGuard is a new tool, there are already many examples of its utilization in the literature that confirm its value [44, 42].

BUTAC.IT AND BUFALE.NET: Are two fact-checking websites that maintain public available black lists of Italian social media pages and websites sharing fake news or misinformation. Some of the categories reported are: pseudo-journalism, conspiracy theories and fake news. Both sites are widely used in studies and analyses in the Italian environment by governmental agencies [7] and researchers [40].

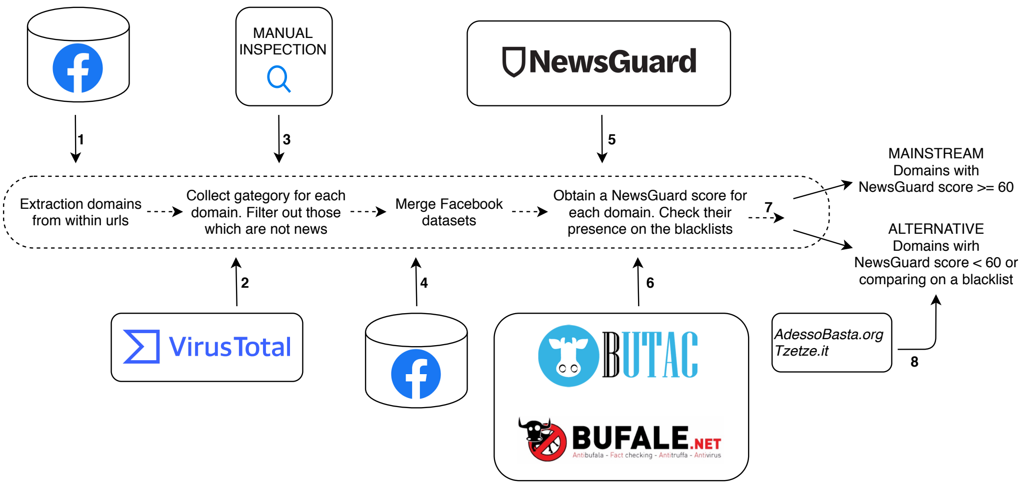

Figure 3.1: Visual representation of the classification pipeline used for labelling the news outlet domains into mainstream or alternative.

3.2.4 *Classification pipeline*

In Figure 3.1 we show the steps faced for the classification process.

Step 1. We extract and collect the domains contained in all the URLs occurring in the *Facebook* dataset. In total we have 638 different domains.
Step 2. We use Virus Total to assign to each domain one or more categories obtained from the report. We then label as "non news" the domains in which none of the following categories appear: news and media, news, blogs, streaming media, blogs, personal sites.
Step 3. Doing a manual inspection we find out that few less-known domains have not received a category by Virus Total. We manually check each of them and label as "non news" domains that are clearly not sharing news.
Step 4. We merge together all labelled domains coming from Step 3 with the domains contained in the *Facebook general* dataset. All the domains from the latter dataset are already considered as news for the way the dataset is created. See Section 4 for more details.
Step 5. Using NewsGuard, we look for a score for each domain. Being this tool new in Italy, and being most of the domains from Italian websites, we are able to extract a score for just 241 domains.
Step 6. We check all our domains on *Bufale.net* and *Butac.it* and label with the assigned category those occurring on them.
Step 7. We classify as mainstream domains with a NewsGuard score above or equal to 60. We classify as alternative all domains with a NewsGuard score below 60 or appearing on at least one blacklist.
Step 8. Finally, we also classify as alternative two more domains: *AdessoBasta.org* and *Tzetze.it*, which have been the focus of several investigations by fact-checkers[7] and debunkers[8], but somehow have escaped our classification process.
The results of our classification process are shown at Table 3.1. Out of 764 different domains, coming from both Facebook datasets, we are able to classify 292 of them (38%). If we consider only the domains appearing in our *Facebook* dataset, we classify successfully 206 domains out of 379 (54%).

Table 3.1: Statistics regarding the domains classified through the classification process.

Domain Category	Count
Total	764
mainstream	222
alternative	60

7 https://urly.it/35dh-
8 https://urly.it/35dhx

3.2.4.1 *Case Study*

To prove the reliability of the pipeline used, we believe it is important to show few concrete examples of the results obtained from the classification process. In Table 3.2 it is possible to see the 5 domains which got the lowest NewsGuard's score. At the top of the table there is *jedanews.it* which, like most of them, appears also on both the black lists - Butac.it, Bufale.net- and is labelled by them as pseudo-journalism. With the third lowest score there is *sputniknews.com*, which is a news websites established by the Russian government, that recently has been accused by Nato of distributing misinformation [9].
Well known established news outlet, like *ilsole24ore.it*, are classified as mainstream and received full score by NewsGuard.
We also made an effort to find few examples of verified fake news, shared in our dataset, later debunked by Butac.it or Bufale.net. Here are the titles of two news which URLs appear in our dataset: 1)'Migranti, eritreo di 22 anni muore di fame dopo lo sbarco'[9], 2)'Salvini toglie finalmente la scorta a Saviano'[10]. The two fake news were shared respectively by *silenziefalsita.it* and *notizieinmovimentonews.blogspot.com*. Both news have been later debunked and pointed out as false ([11],[12]). Both domains are classified as alternative by our pipeline.

Table 3.2: Top 5 domains with the lowest NewsGuard score.

Domain	NewsGuard score
jedanews.it	5.0
ilprimatonazionale.it	5.0
informarexresistere.fr	7.5
ilpopulista.it	12.5
sputniknews.com	12.5

3.2.5 *Temporal Analysis*

In this part of the research, we explore the temporal dynamics related to the dissemination of news, implemented by the accounts studied. Using data analysis and visualization techniques [15], we look at the differences of behavior in posting content from mainstream news outlets and alternative ones. We analyse the patterns used for the two categories of news and compare them with the ones found on Twitter. We look at when the accounts made their first appearance during the time span, the volume of news pushed every week, the hours with

9 https://urly.it/35dh_
10 https://urly.it/35dj1
11 https://urly.it/35dj3
12 https://urly.it/35dj4

more activity, and how the URLs get consumed on the same platform. The ultimate goal is to find differences in the trends implemented while posting alternative and mainstream news showing that the two types of content are pushed following specific rules.

3.2.6 *Content Analysis*

Keeping our focus on how accounts share news, we study the text content of the posts that contain a news URL.
Using natural language processing techniques [23], we see the polarization of posts and the sentiment in relation to the most important entities of the 2018 Italian elections. Finally, we see how the use of hashtags varies in the two platforms and in relation to the two types of content. As for the previous analysis, also during the inspection of the content we mainly look at the differences for posts containing alternative news and those containing mainstream news.

3.2.7 *Influence Analysis*

Since one of the goal of malicious users is to manipulate the public opinion of other users and expand the wave of misinformation they share, we analyse the influence that content pushed by Facebook accounts have on Twitter.
In order to conduct this analysis, we apply a statistical approach called Hawkes model [34]. The idea is to model Facebook and Twitter as collections of point processes. When an event occurs on one of the platforms - an URL being shared - it causes an impulse response on the others platforms as well as on itself, incrementing the probability of a new event related to it to happen - the re-posting of the URL. For each URL, we are able to compute the influence that a specific platform exerts on its occurrences. Summing up the values, we obtain a score representing how much a platform influence the other.
Using this approach, we are able to understand the influence that Facebook have on Twitter and vice versa. The influence score represents how likely a news posted on a specific platform cause the re-share, of the same URL, on the other platform. This results measure the effectiveness of Facebook campaigns on another platform. In order to have a baseline of general Facebook behavior to compare the results with, we repeat the experiments using the *Facebook general* dataset.

4 DATASETS DESCRIPTION AND GENERAL CHARACTERIZATION

In this section, we explain in greater detail the data on which the research is carried out. Then, we show a preliminary exploration of the datasets aimed at identifying their most relevant characteristics.

4.1 DATASETS

For each dataset, we give an explanation about how it is collected and what it contains. Their main statistics are summarized at Table 4.1 and Table 4.2.

4.1.1 *Facebook*

Facebook dataset is built by collecting posts from a set of 23 Italian accounts, using a Selenium crawler [1]. The accounts studied are chosen for the quantity of disinformation they have spread, which was reported to us by a journalist experienced in fact checking. The crawler collected daily posts, for almost a year, for a maximum of 1000 posts per account. The posts are mainly distributed over the period preceding, and immediately following, the 2018 Italian elections.
The data format, returned by the crawler, is to be considered dirty and it requires some pre-processing steps to clean it. The initial data, in fact, is presented in the form of just four attributes: *Date*, *Event*, *Url* and *User*.

- *Date* contains the time when the post was created. When this information is not available, the crawler inserts the current time.
- *Event* represents the whole post in a text format. From it, we extract the text content (corpus), the information regarding the origin, and the type of content, of the post. The origin, when present, is a name representing the account from which the content is shared. The type of content, instead, tell us if the post contains only textual data or also media.
- *Url*, when present, represents the link to an external resource such as a news article. This link is returned in a shortened way [2].
- *User* contains the name of the account which created the post.

1 https://www.selenium.dev
2 https://urly.it/364k6

Table 4.1: *Facebook*, *Facebook general*, and Twitter insights about unique URLs.

Platform	Unique urls	#Alt.	#Main.	Ratio Alt./Main.
Facebook	15822	12744	1653	7.71
Twitter	87970	9231	49633	0.18
Facebook g.	52144	4236	32279	0.13

To this initial features we apply the following pre-processing steps:

1. Conversion of URLs from the compressed to the extended form, using a python library called *BeautifulSoup* [33].

2. After the classification of news domains, we filter out those posts that do not contain an URL. Here, we also remove those posts containing an URL from one of the domains classified as "non news" (see the previous Section).

3. Using two python libraries: *Netpeak* [29] and *3knewpaper* [6], we extract some important information from about URL and the article related to it, such as: status code, title, description, keyword, text author and publish date.

The final dataset is composed by 31 453 posts containing 15 822 unique URLs, posted by 23 different accounts. Out of these 23 accounts, 6 are users, 13 are pages, and 4 are now closed, therefore this information cannot be found. One of the 4 closed accounts was actually shutted down by Facebook itself as a counter-action aimed at eliminating a group of pages, known for sharing misinformation, reported by Avaaz [1].

4.1.2 *Twitter*

For Twitter we deploy an infrastructure which collects the 1% of daily available tweets using the Twitter Streaming API. From the stored tweets, we looked for those containing an URL, posted during the Italian elections period, that matched the keywords and domains contained in the posts of our *Facebook* dataset. More information about the crawling procedure are mentioned in the previous chapter. The fact that Twitter provides public APIs for data collection has made easier to create a clean and complete dataset. After the classification of news domains, we filter out those tweets containing an URL labelled as "non news". The final dataset has 115 497 posts containing 86127 unique URLs shared by about 41k unique users.

Table 4.2: *Facebook* and Twitter insights about posts.

Platform	#Posts	%Alt.	%Main.	Ratio Alt./Main.
Facebook	31453	83.47	8.19	10.18
Twitter	115497	11.25	57.86	0.19

4.1.3 *Facebook general*

Facebook general is used only during the influence estimation experiments. The dataset is publicly available [3]. It consists of observations of the Facebook engagement around Italian political news from *01/09/2017* to *04/03/2018*. In the original dataset there are 84 815 unique URLs, with their respective volume of Facebook interactions (*reactions, comments,* and *shares*) observed every two hour for a week past publication. *Facebook general* structure is very different from the other two previous datasets. Instead of users posts, here, rows represent the interactions gathered by a specific URL in a time interval of 2 hours. There is no information regarding the accounts which shared the URL, or about the text content of the posts. This differences make it impossible to use *Facebook general* in all our analysis, but it fits well our influence studies. For this reason, we use it as a baseline of general news posting behavior on Facebook.
Due to its different structure, the use of this datasets require the following pre-processing steps:

1. We filter out those URLs labelled as "non news" during our classification process.

2. For each row, using the number n of "shares" gathered by the URL in the specific time bin, we unwind the dataset into n singular events. In this way we obtain a structure where each row represents an URL occurrence at a specific time. All the events unwinded by the same row are considered to be happened at the same time. This procedure create more than 4M events.

3. We reduce its size by sampling 1% of the URL occurrences. This sampling rate is the same actuated for Twitter. This decision is taken in order to having it matching the same order of magnitude as the Twitter dataset.

The final dataset contains 52144 unique URLs.

3 https://urly.it/360h2

4.2 GENERAL CHARACTERIZATION

In this section we aim to perform a first exploration of the data. For this purpose, we focus out attention at the main insights of alternative and mainstream news contained in our datasets.

4.2.1 *Mainstream and alternative ratio*

To get a sense of the content shared by *Facebook* and Twitter datasets, we start by retrieving the number of posts and unique URLs, classified either as mainstream or as alternative.
Looking at Table 4.1, the first thing we notice in Facebook is that the number of unique alternative URLs is much higher than the mainstream one. In fact, it is almost seven times bigger than it. Looking at the number of alternative posts, shown in Table 4.2, the trend is similar and the ratio of alternative posts over mainstream ones appears even higher, arriving at 10.18.
The same attitude, instead, is not shared by Twitter. Here, both the numbers of mainstream unique URLs and posts are bigger than the number of alternative ones, and the ratio stays for both URLs and posts below 0.2 .
Taking into account that the number of domains we have classified as mainstream is much greater than the number of alternative ones (see Table 3.1) and considering Twitter an example of a more balanced environment, Facebook sharing behavior is unusual.
As a matter of fact, the quantity of alternative URLs shared by such a small set of accounts, cannot be underestimated and considered as normal. This give us a first prove that the Facebook accounts analyzed share a large quantity of dubious content. Furthermore, they not only share more news coming from alternative domains, they also tend to re-post them far more than for mainstream domains.

4.2.2 *Popular domains*

We now study the popularity of the news websites on the two platforms. In order to do that, we count the number of occurrences for each domain and show the top 10 popular ones for both alternative and mainstream. As additional information, we use a '*' to mark those popular domains occurring in the top 10 for both Facebook and Twitter. If we look at Table 4.3, we can see that for Facebook a very high number (95%) of alternative posts is focused on just the top 5 domains. Among this group, there are some of the recently investigated[4,5] news outlets, such as as *adessobasta.org*, *silenziefalsita.it*, and *notizieinmovimentonews.blogspot.com*, accused of sharing misinformation. Furthermore,

4 https://urly.it/35dhx
5 https://urly.it/35dht

Table 4.3: Top 10 popular domains on Facebook and respective percentages of events.

Alt. Domain	**%Alt.**	**Main. Domain**	**%Main**
adessobasta.org	29.69	*ilfattoquotidiano.it**	27.77
silenziefalsita.it	18.42	*beppegrillo.it**	10.47
direttanfo.blogspot.com	17.36	*repubblica.it**	6.90
lonesto.it	15.40	*corriere.it**	6.09
notizieinmovimentonews.blogspot.com	14.11	*huffingtonpost.it**	5.86
*ilblogdellestelle.it**	3.41	*ansa.it**	5.82
*liberoquotidiano.it**	0.18	*fanpage.it*	4.23
*infosannio.wordpress.com**	0.14	*ilgiornale.it**	3.72
*dagospia.com**	0.13	*lastampa.it**	2.75
*sputniknews.com**	0.12	*ilmattino.it*	2.72

both the debunked fake news we that have found, reported in chapter 3, belong to this domains. For mainstream, instead, even if there is an high number of posts coming from *ilfattoquotidiano.it*, the overall trend is smother and more distributed over more domains.
On Twitter (see Table 4.4), we can see a similar behavior with an higher concentration in less domains for the alternative category compared to the mainstream one, but not as sharp as for Facebook.
This results show that misinformation on social platforms is pushed from few specific domains.
Another interesting finding is that while for the first top 5 Facebook alternative domains there is no counterpart on the Twitter table, the mainstream ones have it. On Twitter the same phenomenon is happening for *voxnews.info*. In fact, the Twitter top alternative domain, appearing in 27% of the tweets gathered, does not appear on the Facebook popular ones.
In general, while 80% of mainstream domains are popular on both Facebook and Twitter, only 50% of alternative ones are so. Then, it appears that popular alternative domains are quite unique for each platform. Therefore, being used on a platform does not guarantee the popularity on the other.

4.2.3 *URLs occurrences*

Next, we investigate Facebook and Twitter posting behavior studying how many times each news appears on a specific platform. For both Facebook and Twitter we plot the CDF of the occurrences of each unique URL. We observe that on Twitter (see Figure 4.2) more than 80% of URLs appear only once and few URLs are posted more than

Table 4.4: Top 10 popular domains one Twitter and respective percentages of events.

Alt. Domain	%Alt.	Main. Domain	%Main
voxnews.info	19.00	*repubblica.it**	15.68
*ilblogdellestelle.it**	13.09	*ilfattoquotidiano.it**	7.84
*liberoquotidiano.it**	10.68	*corriere.it**	7.46
*dagospia.com**	8.28	*lastampa.it**	7.13
scenarieconomici.it	8.18	*beppegrillo.it**	5.33
imolaoggi.it	4.79	*ansa.it**	5.32
ilprimatonazionale.it	4.49	*huffingtonpost.it**	3.61
*infosannio.wordpress.com**	3.17	*ilgiornale.it**	3.51
globalist.it	2.82	*ilsole24ore.com*	3.36
*sputniknews.com**	2.45	*linkiesta.it*	2.10

100 times, but overall the distribution is very similar for mainstream and alternative news. On Facebook, instead (see Figure 4.1) more than 40% of URLs appear only once, but there is a substantial difference between alternative and mainstream. Alternative URLs, in fact, tends to appear at least twice for more than 60% of the times, and very few URLs are pushed more than mainstream ones.
Considering that on Facebook we are studying just 23 accounts, a news shared more than one time, for us, could mean mainly two things: 1) The same account is pushing the same news multiple times. 2) One or more accounts act as an echo of the source account in order to extend the visibility of the content. We study this behavior more in details in the next chapters.

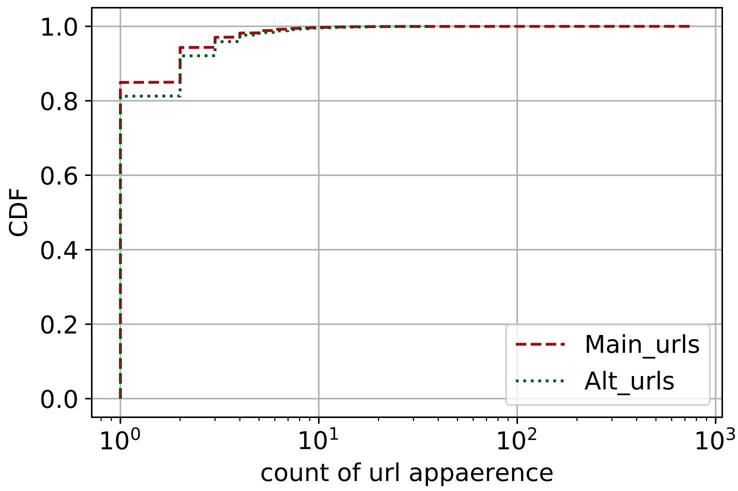

Figure 4.1: CDF of the URLs occurrences on Facebook

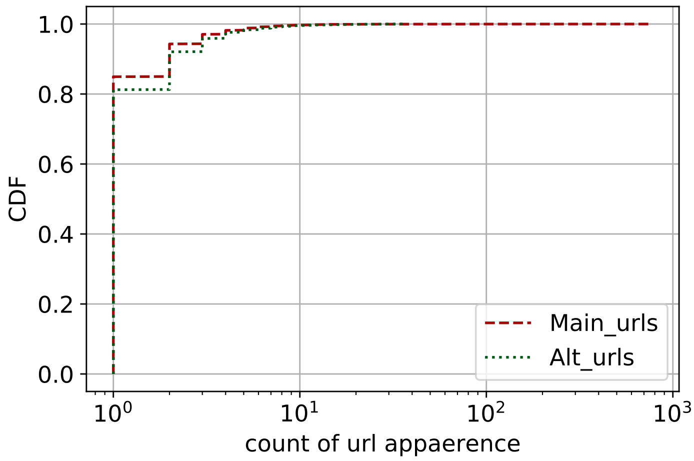

Figure 4.2: CDF of the URLs occurrences on Twitter.

4.2.4 *Take-Aways*

Summarizing, our general characterization shows the following findings:

1. The quantity of alternative news spread on Facebook by a limited set of accounts, and the high volume of URLs coming from few proved low-quality domains, confirms the occurrence of disinformative campaigns in the 2018 Italian panorama.

2. For both Facebook and Twitter, alternative news are much more focused in few specific domains than mainstream ones. Moreover, the most popular domains are very specific for each platform.

5

TEMPORAL DYNAMICS

In this section, we study how the phenomenon of news dissemination occurred over time. Our goal is to find the patterns followed in the process of sharing mainstream and alternative news.
For the purpose of exploring the temporal behavior of our *Facebook* dataset, we filter out around 9k rows without a reliable date. This process produce Table 5.1, which is the one we use for the following analysis. Moreover, to better understand the underlying distribution of the data, we take as reference points two important dates for the 2018 Italian elections: The election day - 04/03/2018 - and the day of sworn for Conte Cabinet - 01/06/2018.

Table 5.1: Facebook filtered dataset 1.

#Posts	%Alt.	%Main.	Ratio A./M.
23627	81.02	10.43	7.76
Unique URLs	**#Alt.**	**#Main.**	**Ratio A./M.**
14424	11368	1642	6.92

5.1 FIRST ACTIVITY AND PERCENTAGE OF ACTIVE USERS

We now look at the point in time when the Facebook accounts made their first activity in our dataset. Then, we compute for each week, in the time frame measured, the number of accounts which posted their first news in that week. Looking at Figure 5.1, it's possible to see that, as expected, all accounts started their activity before the election

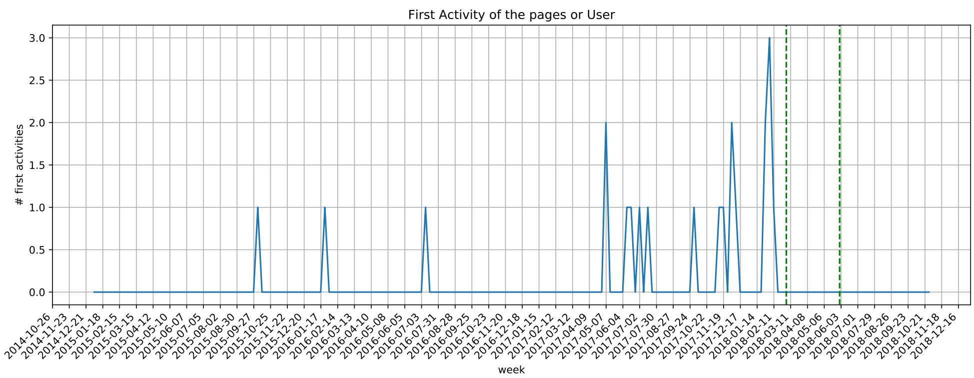

Figure 5.1: Counts of the first activities of the Facebook accounts, grouped by week.

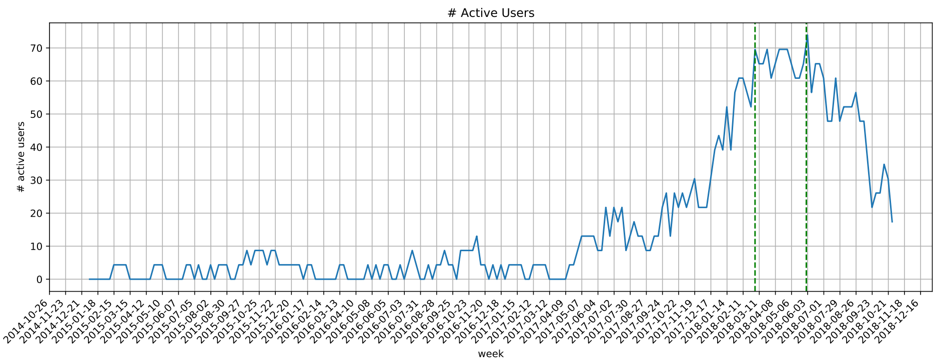

Figure 5.2: Percentage of active Facebook users, grouped by week.

day. Looking at the graph, it is interesting to notice that there is a group of pages activating very close to the elections. In fact, 12 pages out of the 23 total, made their first appearance between October 2017 and February 2018. Interestingly, 6 pages of this group did their first activity immediately after January 28, which was the day in which the electoral lists were presented.

It seems than many of this accounts were triggered by the approaching of the elections or by events related to them. As a result, they started posting new content just for the election period, even if most of the accounts were created a while before. This result make us wonder if all the accounts stay active through the complete time span or if they activated only for specific occasions. For further investigate this behavior, we now plot the percentage of unique active accounts per week.

From Figure 5.2, it can be seen that the percentage of active unique accounts is constantly increasing approaching the elections. From the election day to the day of sworn for Conte Cabinet there is a general high trend which reaches its peak exactly in the week of the latter event. Then activity starts to decrease again.

Our idea it is that most of the accounts were created from scratch, or re-activated on purpose, to push their content during the election period. Rather than influence the election outcome, the accounts seems to focus their efforts between the election day and the official starting day of the new government, in order to manipulate the public opinion related to it.

5.2 WEEKLY URLS OCCURRENCE BEHAVIOR

We next explore the quantity and the distribution over time of news shared. We start by computing the number of shared URLs grouped per week, normalizing the results by the average number of weekly posts made on each platform. From Figure 5.3, it is possible to see that the sharing behavior on Facebook follow very much real word

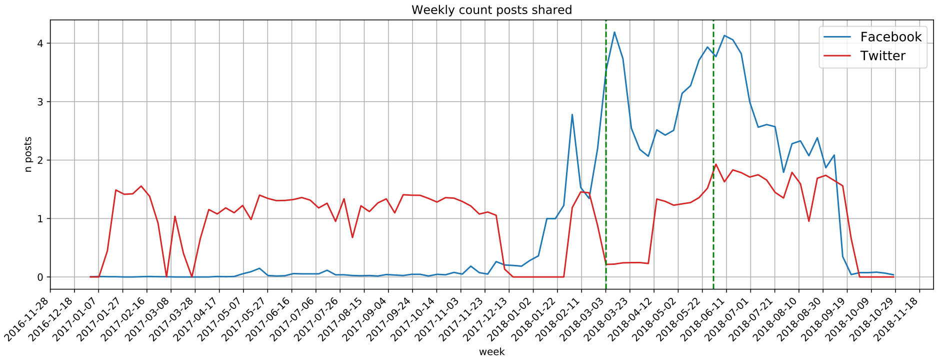

Figure 5.3: Facebook and Twitter normalized number of events per week.

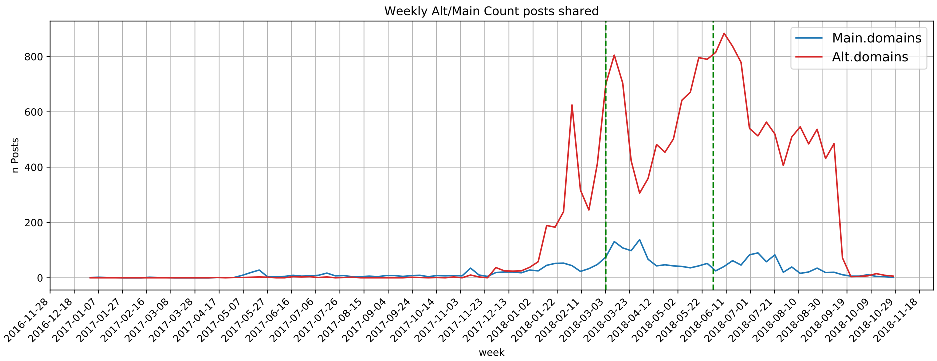

Figure 5.4: Weekly number of alternative news URLs and Mainstream news URL shared on Facebook.

events. The distribution has two big spikes which occur close to the two most important events of the Italian election. The spike that happens around February 2018 seems to be caused by an error in our crawler. Twitter, instead, is much more equally distributed over time, even if its distribution too has its highest point coinciding with the day of sworn for Conte Cabinet. Interestingly, despite the fact that the Twitter dataset is crawled from 1% of the whole Twitter, there are two lows in the distribution. The first one happens on January 2018 and, considering that there is not even a post, it represents an error in our crawling infrastructure. The second one, instead, happens on March 2018 right after the election day, and it goes on for a couple of weeks. This low is not an error because actually some data are retrieved. What could it be then? We think that it could be related to the electoral silence requested for the day of the election and the ones preceding it. The fact the the low is more visible after the elections and not before them could be due to the choice of grouping by week. It is interesting to find out that, while the electoral silence is somehow visible on Twitter, we could not see sign of it on Facebook.
Let's now focus our attention on the differences applied for alternative news and mainstream ones on each platform. Looking now at Figure 5.4, we can see that, on Facebook, for the whole period, the

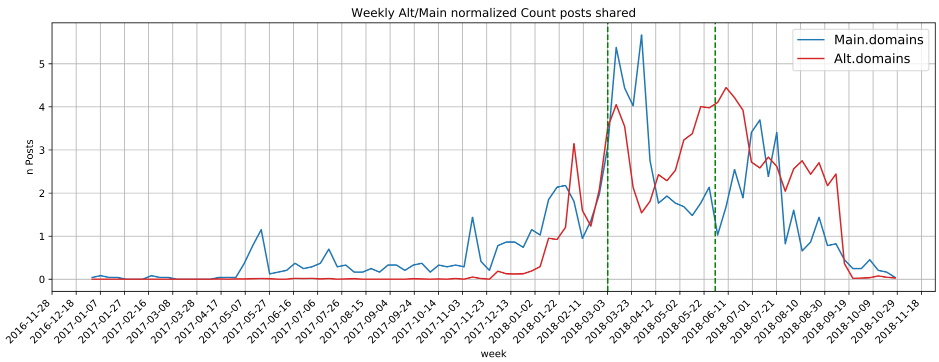

Figure 5.5: Weekly number of alternative news URLs and Mainstream news URL, normalized by the weekly average number of shared URLs for each category, shared on Facebook.

number of alternative posts outperforms the number of mainstream ones. For Twitter it is the opposite (see Figure 5.6). We have already seen this trend in the general characterizations of the two datasets but, from the second Figure 5.6 it is possible to notice a new detail. On Twitter, the only time period in which alternative posts exceed the mainstream ones is in the already cited period immediately after the election day. This result reinforce our idea that on Twitter, given the bigger pool of users present in the dataset, it can be seen the effect of the electoral silence requested on social media [10], and that mainstream content is more prone to respect it.
Finally, we look at the number of alternative and mainstream posts normalized by the average weekly number of posts made by that category over that platform. We want to find out when each type of content was pushed more than usual on each platform.
Starting from Facebook (see Figure 5.5) we can observe that alternative news are shared more, close to important events. In fact, the distribution follow very well their occurrences. Mainstream news, instead, are pushed more after events, that is when their distribution reaches its peaks. This trend could be summarize with the assumption that while alternative news are pushed both before and after events, with the aim to manipulate the public opinion, mainstream ones are shared later, to inform about the event itself. On Twitter (see Figure 5.7) we note that as soon as the election period starts, even if the total number of posts is much smaller, alternative content is pushed more than on average with respect to mainstream. This result represents an unique behavior that might mean that the large quantity of Facebook content is arriving on Twitter too, influencing it. Or simply that generally during the elections period alternative content is posted more aggressively.

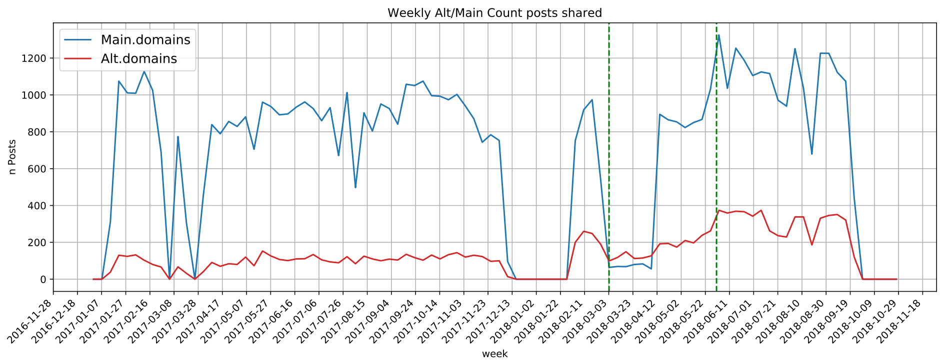

Figure 5.6: Weekly number of alternative news URLs and Mainstream news URL, shared on Twitter.

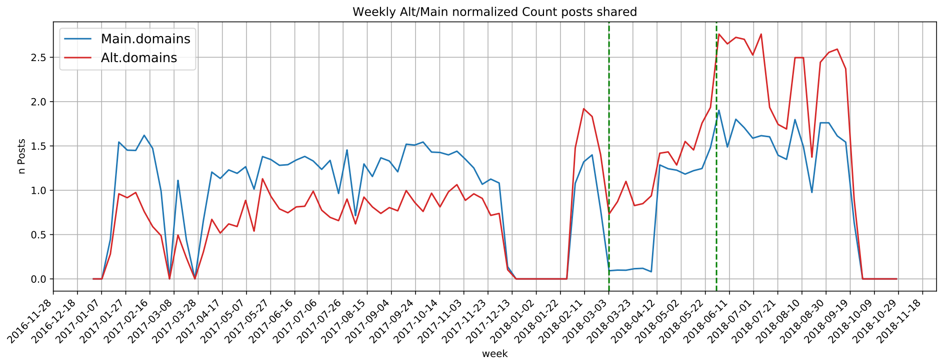

Figure 5.7: Weekly number of alternative news URLs and Mainstream news URL, normalized by the weekly average number of shared URLs for each category, shared on Twitter.

5.3 HOURS OF THE DAY AND HOURS OF THE WEEK

Next, we look at the hours of the day and week when posts are mostly shared. To do this kind of analysis, as well as all the future ones that need the exact posting time, we have to filter out again our *Facebook* dataset. Starting from Table 5.1, we remove around 3k rows without an exact time, producing Table 5.2.

Table 5.2: Facebook filtered dataset 2.

#Posts	%Alt.	%Main.	Ratio A./M.
20079	81.08	11.27	7.19
Unique URLs	**#Alt.**	**#Main.**	**Ratio A./M.**
13.14	11033	1538	7.17

To conduct this analysis, we first explore the hourly distribution of posts, checking the percentages at each hour of the day. Note that all the hours in the graphs are reported in UTC. From Figure 5.8, it is possible to see that Facebook distribution is very fragmented but it

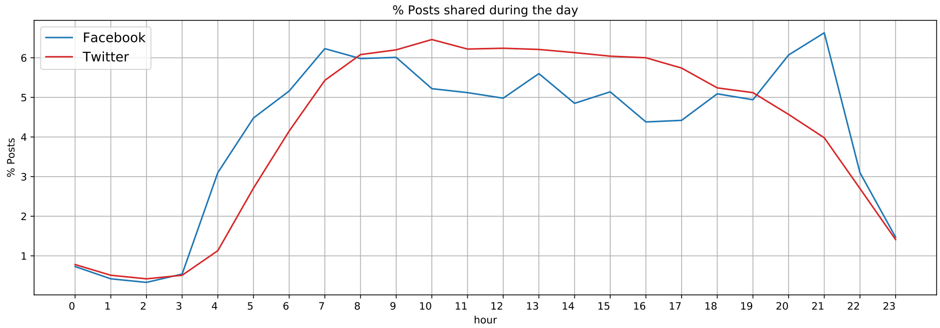

Figure 5.8: Facebook and Twitter percentage of shared events per hours of the day.

has most of its peaks at "sides hours", very early in the morning or quite late in the afternoon. This is different from Twitter which activity is constantly increasing until noon, when it starts to slowly decrease. We decide to investigate each platform individually, comparing the hours of the posts containing a mainstream URL with the ones with an alternative one.
From Figure 5.9, we can see that on Facebook the previous behavior is mostly caused by alternative news, which, as a matter of fact, are highly shared at "sides hours" of the day. Alternative posts are mainly created at 9 pm or at 7 am while mainstream are more equally distributed in the afternoon, with an increasing trend that starts at 12 pm and reaches its peak at 6 pm. Without leaving Facebook, we now look at the data plotted considering the hours of the week, see Figure 5.10. The difference of posting behavior for the two categories is even more visible now. In fact, the two distributions almost never overlap. Mainstream has its highs on Wednesday and Friday in the mid afternoon, while alternative does not have preferred days but is more distributed over the week.
Twitter, instead, is generally more regular. This uniformity is given by the high number of different users gathered in the dataset. Replicating the analysis done for Facebook, we can see in Figure 5.11 that from 2

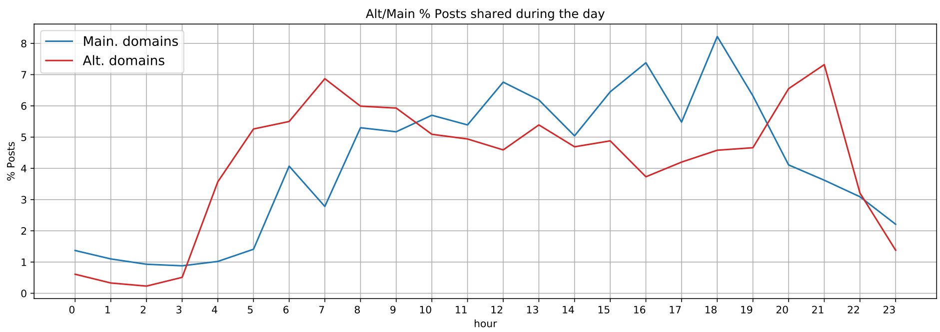

Figure 5.9: Alternative and mainstream percentage of events shared per hours of the day on Facebook.

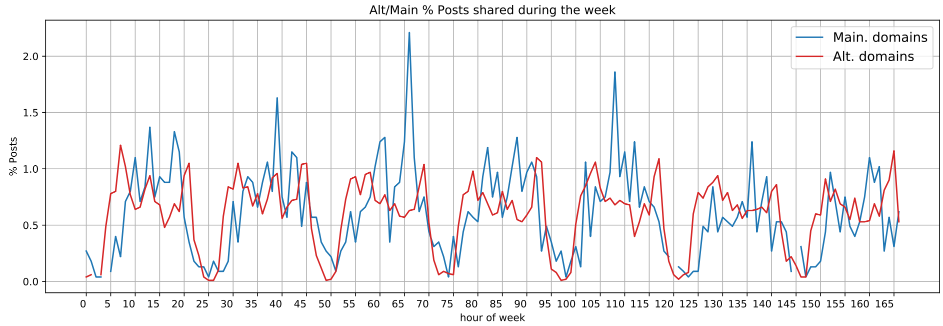

Figure 5.10: Alternative and mainstream percentage of events shared per hours of the week on Facebook.

pm onward alternative URLs are more shared than mainstream ones. Interestingly, there is an high volume of alternative news shared at 7 pm. Comparing the two distributions with some background data about the popular posting hours on Twitter [19], we find out that while mainstream distribution matches the usual Twitter behavior, alternative news are shared at different times. For the hours of the week (see Figure 5.12), it visible that, especially on some particular days like Wednesday and Sunday, the two categories of news are shared at different hours.

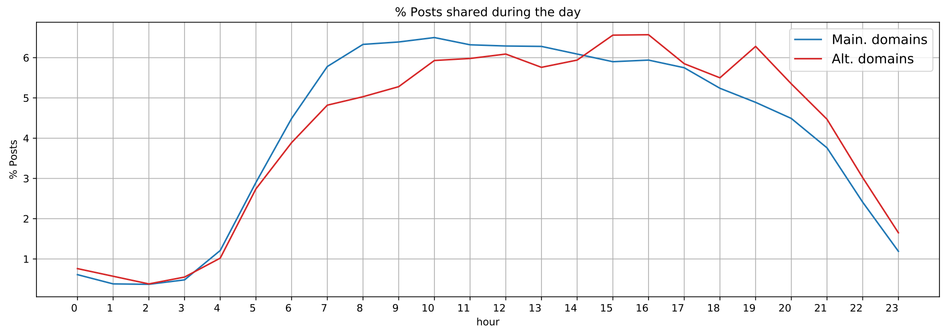

Figure 5.11: Alternative and mainstream percentage of events shared per hours of the day on Twitter.

5.4 URLS CONSUMPTION

We now look at URLs that appear more than one time in our dataset, in order to understand how news get consumed on social platforms. We start by plotting the CDF of the difference between the first occurrence of an URL and its next occurrences within the same platform (see Figure 5.13). First of all, for both Facebook and Twitter, few news are recycled even after a long time span - E.g., 1000 hours (41 days). Then, it is possible to see how platforms show different trends with respect to the categories of content. While on Twitter alternative URLs

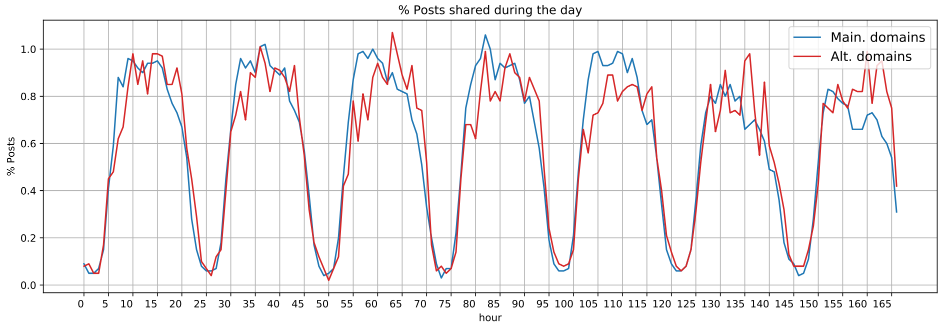

Figure 5.12: Alternative and mainstream percentage of events shared per hours of the week on Twitter.

spread faster than mainstream ones, on Facebook it happens the opposite. In fact, on Facebook 50% of mainstream URLs compared to 25% of alternative ones, are re-posted after their first appearance, in less than 0.1 hours ~6 minutes.
To assess the statistical difference between the two distributions we perform a two-sample Kolmogorov-Smirnov test [24]. This kind of test is used to check whether two samples come from the same distribution. If the *p-value* returned is above a threshold, usually 0.001, then the test is said to confirm the null hypothesis, supporting that both samples come from a population with the same distribution.
In our case, the test rejects the null hypothesis between the distribution of mainstream news and alternative news ($p<0.001$), confirming that the two category of the content come from different distribution. This means that each category influence how the URL is consumed.
Looking again at both platforms together, it is clear that on Facebook both alternative and mainstream news spread much faster than on Twitter. This behavior can be caused by the specific structure of the accounts that we are analysing on Facebook. It is possible that one or more accounts act as an echo of bigger pages to increase the visibility of their content. This theory would explain the speed with which content is being re-shared over the platform.
Next, we look at URLs that appear on both Facebook and Twitter. We plot the CDF of the time difference between the first occurrence on a platform and the next occurrence on the other platform (see Figure 5.14). It is interesting to notice that weather the news starts from Facebook and later appears on Twitter, or starts from Twitter and moves to Facebook, alternative news have a smaller time gap between the two occurrences. This result is consistent with previous study [44] on the dissemination of news, executed on different platforms. This means that it represents a general pattern for this kind of content.

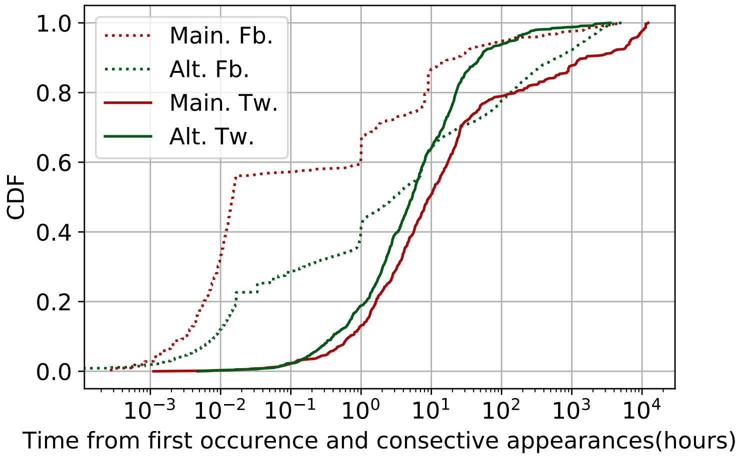

Figure 5.13: CDF of time difference between the first and the consecutive occurrences within the same platform.

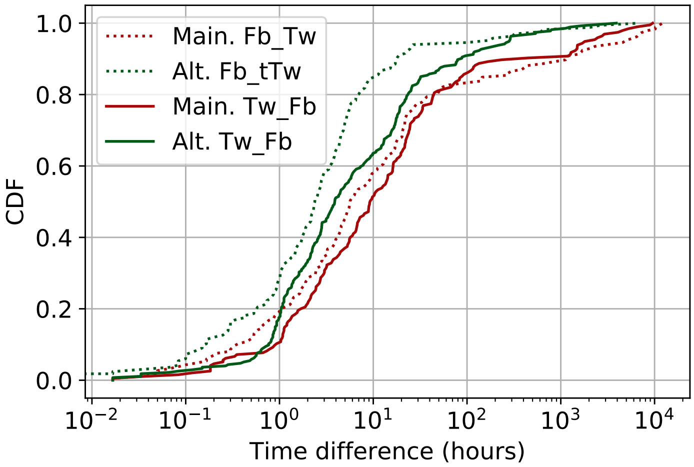

Figure 5.14: Time difference between the first occurrence of a URL on a platform and the following occurrence of the same URL on the other platform.

5.5 ECHO-STRUCTURE

In the previous analysis we noticed that many URLs in our Facebook dataset are re-shared within a small amount of time. This behaviour is evident in Figure 5.13, where, for Facebook, both categories of news propagate faster than on Twitter. The difference between the platforms is that evident that we decide to further investigate this matter. Going deeper through our analysis we find out an interesting behavior implemented by some pages in our *Facebook* dataset. Many of the URLs shared by more than an account are actually shared always by the same accounts. This raise an interesting question: Are some accounts acting like an echo of bigger pages ?
To address this question, we compute for each couple of accounts

the percentage of matching URLs. Not surprisingly, we find that some pages share content mostly originated from other accounts. For example, to keep the account's names anonymous, let's call two pages in our dataset A and B. 80% of the URLs shared by page B are shared also by page A. If we now look at the temporal dynamics, we can see that 95% of the URLs shared by both the pages are shared first by A and then by B. Finally, comparing the number of followers, we see that A is much bigger than B. This is a proven evidence that B acts only as an echo of A, sharing most of the latter content in order to increase its visibility. This is just the most striking example of this kind of behavior visible in our data, but it is not the only one. Out of the 23 accounts analyzed, we find 3 pages (13%) that have at least 40% their content originated from other bigger pages.

5.6 TAKE-AWAYS

In summary, the temporal dynamics analysis yields the following findings:

1. Most of the accounts in the Facebook dataset were triggered by the approaching of the elections and became very active in pushing alternative news just for that specified period.

2. On Facebook, the spikes of the quantity of content shared occur close to important events, showing that the implemented campaigns are actually motivated by real world events.

3. On Facebook, while alternative news are heavily pushed both before and after events, mainstream ones are mainly shared after.

4. For both the platforms, there are different times in which high quantities of mainstream and alternative news are pushed, with the latter category predominant at specific hours.

5. For both the two platforms, the rate of spreading each type of content is different. While on Facebook mainstream news spread faster than alternative news, on Twitter happens the opposite. For news shared by both platforms, alternative URLs have a smaller time gap between the first occurrence on a platform and the following on the other platform.

6. 13% of the studied Facebook accounts act as a echo of bigger pages to increase the visibility of the original content.

6 CONTENT ANALYSIS

In this section we further investigate the content of posts. In particular, we explore the sentiment of the text contained in the posts and the hashtags used. This analysis rely on the complete *Facebook* dataset (see Table 4.1 and see Table 4.2).

6.1 SENTIMENT ANALYSIS

We analyse the emotional dynamics related to the spread of mainstream and alternative news. To do so, we assess the sentiment of each post, for both Facebook and Twitter, using the *Polyglot* library [31]. Sentiment analysis is the interpretation and classification of emotions - positive, negative, and neutral - inside text data, using automated processes [30] . Therefore, for each post, we obtain a score between -1, which represents negative sentiment, and +1 which, instead, represents a positive one. The sentiment is intended to show the emotional attitude of Facebook or Twitter users when posting that content. Setting a threshold equal to 0.01 to the sentiment score, we are now able to classify posts into three categories: Positive [0.01, 1], Neutral (-0.01, +0.01), and Negative[-1, -0.01]. The threshold is obtained after attempts around values used in previous work [18].
Table 6.1 contains the results obtained. It displays the percentages of scores for each sentiment, for alternative and mainstream news, computed for Facebook and Twitter. Firstly, looking at the differences between the two platforms, it can be seen that overall the studied Facebook accounts show a more positive attitude than the Twitter users. Focusing now on one platform at a time, we see that for Facebook the percentage of posts with a positive score is greater than the percentages of negative and neutral posts. From the ratio, computed diving the number of positive posts over the number of negative ones, it can be seen that as the number of posts considered increases, the ratio increases too. Indeed, the lowest ratio belongs to mainstream URLs, that are just a small part of the *Facebook* dataset.
To assess the statistical difference between mainstream and alternative score's distribution, which we plotted at Figure 6.1 and Figure 6.2 , we apply the two-sample Kolmogorov-Smirnov test. This statistical test is extremely useful because is effective even if the distributions differ in size [24]. The result of the test confirms that the alternative and mainstream score's distributions exhibit a unique behavior, returning a p-value below 0.01. Thus, it is correct to say that the category of news do influence the sentiment.

Table 6.1: Facebook and Twitter results of the sentiment analysis. The Table reports the percentage of posts for each sentiment and the ratio of positive posts over negative ones.

Facebook				
	%Positive	**%Negative**	**%Neutral**	**Ratio P/N**
Total	47.11	37.49	15.35	1.25
Main. news	44.45	40.58	14.96	1.09
Alt. news	45.32	39.34	15.33	1.15
Twitter				
Total	44.27	38.13	17.58	1.16
Main. news	43.06	39.04	17.89	1.10
Alt. news	41.19	40.88	17.91	1.00

On Twitter the results are similar. Overall, there is an higher number of positive tweets than negative and neutral ones. On this platform, alternative news is the category with the lowest ratio. This result confirms the previous theory regarding the correlation between ratio and number of posts. In fact, on Twitter, alternative news represents the smallest fraction of posts. The two-sample Kolmogorov-Smirnov test on the alternative and mainstream score's distributions rejects again the null hypothesis ($p<0.01$), confirming that each category exhibit an unique behaviour also in Twitter.
Finally, we compare our sentiment's scores with the results computed on a general set of tweets, obtained in previous studies [45]. Overall, our sentiment's distributions exhibit a more negative attitude. Our theory is that this behavior could be related to the fact that the posts in our datasets, for both Facebook and Twitter, are strongly related to political matters.

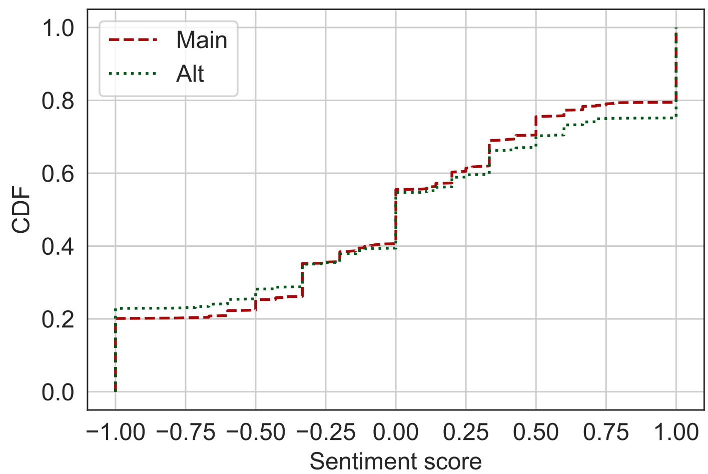

Figure 6.1: CDF of the Facebook sentiment's score distribution

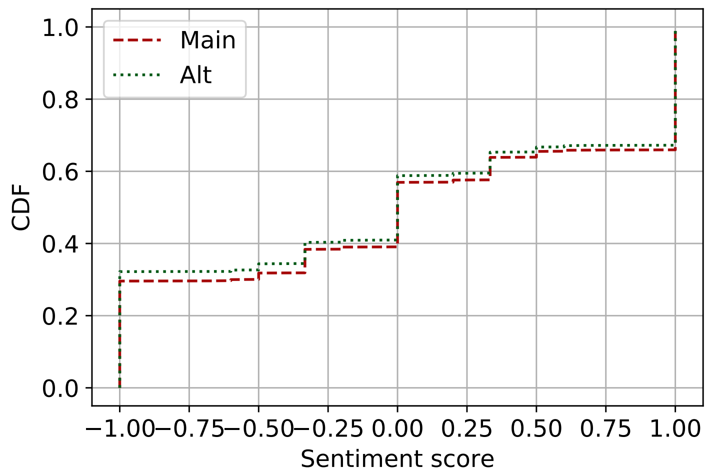

Figure 6.2: CDF of the Twitter sentiment's score distribution

6.1.1 *Entity Sentiment*

To get a deeper understanding of the Italian panorama during the 2018 elections, we plot the distributions of the sentiment's scores related to the politicians of the four most voted parties of that year (see Figure 6.3). For each party we choose its leader, obtaining the following list: Silvio Berlusconi for *Forza Italia (FI)*, Luigi Di Maio for *Movimento 5 stelle (M5S)*, Matteo Renzi for *Partito Democratico (PD)*, and Matteo Salvini for *Lega (LN)* [25].
Starting from Facebook, we find out that Di Maio and Salvini share a much larger portion of posts than the others two political candidates, showing up in respectively 11.7% and 8.4% of all posts. All the score's distributions of the analysed politicians have their median close to zero. This means that the range of sentiments involving posts with their names is various. The only exception is Di Maio. Di Maio sentiment's distribution has its median on average six times larger than the others. Furthermore, very few time its name is associated with a completely negative (-1) posts.
On Twitter,instead, the two most popular deputies are Salvini, showing up in 3% of the all posts, and Renzi which appears in 2% of the tweets. Salvini is the only one with a negative median; however, very close the zero. Di Maio sentiment's distribution, instead, is completely unbalanced towards positive scores, meaning that most of the users using his name share tweets containing positive words.
Overall, the fact that Di Maio seems to have an higher sentiment score while Salvini has a negative one could mean that most of the users in our datasets are supporters of the first one and against the latter. Another possible interpretation for the low score of the posts containing Salvini's name, could be that users supporting him, and therefore using his name, tend to use more negative words.

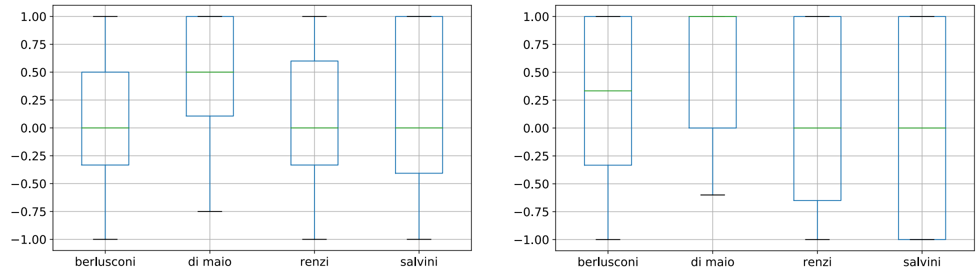

Figure 6.3: Facebook (left) and Twitter (right) sentiment distributions related to each considered politician.

6.2 HASHTAGS ANALYSIS

We assess the use of hashtags in posts. Ultimately, hashtags strongly relate to news content. Therefore, studying the popular ones in our datasets give us an idea of the type of news shared and the message contained in the related post. We first explore the use of hashtags on different social platforms. A study made in 2016 by Buzzsumo shows clearly that, on Facebook, posts with at least one hashtag result in a lower engagement[1]. The opposite happens on Twitter where hashtags have a very important value. In fact, tweets that include hashtags are 33% more likely to be retweeted than those without[2]. Giving those information we are now able to better understand the results of our analysis.

Looking at Table 6.2, we can see that on Facebook just 2% of posts have at least one hashtag, while on Twitter almost 30% of them have it. This percentages are no surprise because they reflect the platforms attitude towards hashtags. Nevertheless, looking at alternative and mainstream news there seems to be differences in the usage of hashtags. Posts that contain an alternative URL are less matched with hashtags. In Facebook this difference is very evident but in Twitter it can be observed too.

To better understand the differences between mainstream and alternative news, we extract the top ten popular hashtags for each platform and type of news (see Table 6.3). On Facebook, most of the hashtags refer to the popular political parties or deputies; among them, for both alternative and mainstream news, we can find *#m5s*, *#lega*, *#renzi*, and

1 https://urly.it/35dyn
2 https://urly.it/35dyp

Table 6.2: Facebook and Twitter stats about hashtags usage.

Platform	% overall hashtags	%Alt.	%Main.
Facebook Pages	2.04	0.62	14.74
Twitter	29.04	19.01	26.32

#salvini. Despite the general low usage of hashtags, we find out that two specific pages tended to accompany their posts with hashtags referring to the name of their account itself. We are talking about *#figlidiputin* and *#nicoiovinee*. Looking for the same hashtags in the Twitter dataset we find no sign of use, meaning that probably those accounts didn't have a counterpart on the Twitter side.
Moving our attention to Twitter, we can see that tweets that share alternative news seem to contain more extremist hashtags in their popular ones. Looking at Table 6.3, while for the mainstream side we can see mostly common hashtags used for news, on the alternative side we can find some unusual ones. For example: *#islam*, *#casapound* which represents an Italian far-right party [12], or *#noiussoli*. No-Ius-Soli represents a slogan used mainly by the right-party against the introduction of a new law that grant citizenship to those born in Italy[3].

Table 6.3: Top 10 popular hashtags, for Mainstream and Alternative news, in Facebook and Twitter.

Facebook				Twitter			
Main.		Alt.		Main.		Alt.	
Hashtag	**(%)**	**Hashtag**	**(%)**	**Hashtag**	**(%)**	**Hashtag**	**(%)**
figlidiputin	0.25	nicoiovinee	0.12	m5s	0.56	m5s	0.08
nicoiovinee	0.19	movimento5stelle	0.05	agi	0.41	ilmiovotoconta	0.05
m5s	0.17	m5s	0.03	roma	0.36	noiussoli	0.04
salvini	0.08	ilmiovotoconta	0.03	salvini	0.35	renzi	0.04
pd	0.07	salvini	0.01	pd	0.29	salvini	0.04
berlusconi	0.06	renzi	0.01	fattoquotidiano	0.27	pd	0.04
renzi	0.05	decretodignità	0.01	renzi	0.27	thexeon	0.03
lega	0.04	lega	0.01	news	0.24	islam	0.03
legaladrona	0.03	governopatrimoniodelpaese	0.01	migranti	0.22	casapound	0.02
putindivista	0.03	mattarella	0.01	raggi	0.18	profughi	0.02

6.3 TAKE-AWAYS

In summary, the content analysis shows the following results:

1. On Facebook and Twitter, overall, the percentage of posts with a positive score is greater than the percentages of negative and neutral posts. However, for both platforms, alternative and mainstream news are associated with different sentiments. A statistical test confirms the difference between the score's distributions of the two types of content.

2. In both platforms, alternative news are less matched with hashtags. A deeper analysis shows that for spreading alternative content on Twitter are often used more extremist hashtags. For

3 https://urly.it/3620c

example, the popularity of hashtags such as #*casapound* and #*noiussoli* might represent the intention of share hate-oriented messages.

7

INFLUENCE ANALYSIS

So far, we have analyzed the temporal dynamics and text content of the posts in our datasets. Alternative and mainstream news have been proven to be shared with different patterns, strictly related to the category of content. However, it is important to keep in mind that often this posts are made for propaganda, trying to influence the public opinion of as many people as possible. Then, it is reasonable to presume that the Facebook accounts activity aims to reach other social platforms. This raise an interesting question: how efficient are the studied accounts at disseminating news towards Twitter?
In this section, we explain how we can model the influence that a platform exert on another, adopting a statistical model called Hawkes processes. First, we explain briefly some basic concepts on point processes, focusing on self-exciting ones. Then, we describe the concepts behind Hawkes processes, and how they can be adapted to fit our situation. Finally, we show the results obtained from the influence analysis, comparing the efficiency of Facebook and Twitter at disseminating URLs.

7.1 INTRODUCTION ON POINT PROCESSES

In statistics and related fields, point processes are collections of points randomly located on some underlying mathematical space such as time and location [6]. Point processes provide us the tools to model the properties and timing of events. For example, figure 7.1 depicts an example of a point process. It represents a cascade of "re-tweets" about a Gaming video where each tweet constitute an event.
Giving a more formal definition, a point process is a random process whose realizations are event times *T1,T2,..Ti* where *Ti* represents the

Figure 7.1: A point process showing the occurrence of tweets about a Gaming video on YouTube. The first 10 events are shown. An event with hollow tip denote a retweet of a previous tweet. Image taken from [34].

time of occurrence of the *i-th* event [34].
The easiest and most studied point process is the Poisson process that consists of stochastic, randomly spaced points [8]. A Poisson process is usually used to model simple problems where the arrival of an event depends only on relevant information about the current time, i.e., it is independent from previous events. Due to this property, a Poisson process is said to be memory-less [38].
To describe more complex problems, such as the arriving of customers during peaks hours, we need be able to vary the event intensity with respect to time. In order to do that, there is a subclass of Poisson processes, called non-homogeneous, in which the rate of events arrivals is a function of time [8].

7.2 SELF EXCITING PROCESSES AND HAWKES

Being able to vary the rate of arrivals with respect to time, is not yet enough to model complex problems like the spread of news on social media platforms.
Let's think about the process of sharing a news URL on Twitter. The arrivals rate of posts containing the same URL depends on time. In fact, there are specific hours at which the volume of tweets is higher. But also the total number of shares and the temporal distance from the creation of the original post influence the following rate of arrivals. As a matter of fact, the number of "re-tweets" will be higher in the first hours after the creation. Furthermore, more shares mean more visibility, which increase the probability of new shares.
For this purpose, we now introduce a new class of processes where the probability of seeing a new event increases due to the presence of previous events. Processes of this class are called self-exciting, and they are characterize by the fact that their current intensity is determined by the past history of the process itself [32]. One well-known self-exciting process is the Hawkes process.
In the Hawkes processes the intensity function $\lambda(t|\mathcal{H}_t)$ depends on the current time t and on the process history $\mathcal{H}_t$. It takes the form:

$$\lambda(t|\mathcal{H}_t) = \lambda_0(t) + \sum_{i:t>T_i} \gamma(t - T_i). \tag{7.1}$$

where:

- $\lambda_0(t) > 0$ is the background rate. It describes the arrivals of events triggered by external sources.
- T_i are the event times occurring prior to time t. The presence of those events as well as their distance $(t - T_i)$ from time t, influence and contribute to the event intensity function at time t.

- $\gamma(.)$ is a function called impulse that expresses the influence of the past events T_i. Typically, the impulse function is taken to be monotonically decreasing to let more recent events having a stronger influence compared to events occurred previous in time [34].
- The summation term explains the dynamics of the "self-exciting" effect.

We observe that the Hawkes processes are a particular case of non-homogeneous Poisson process, in which the intensity is stochastic and it explicitly depends on previous events through the impulse function $\gamma(.)$ [34].

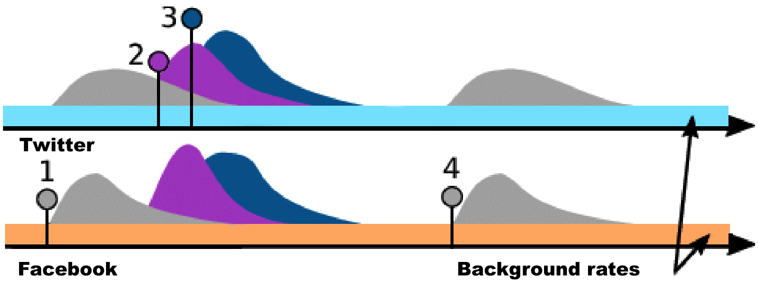

Figure 7.2: An Hawkes model with 2 processes - Facebook and Twitter- in which occurred 4 events.

7.3 HAWKES STATISTIC

Following [43, 44], the idea is to model the diffusion of news on social media platforms fitting an Hawkes model. We treat the analyzed platforms as point processes belonging to the model, each with its own background rate. When an event occurs on one of the platforms, an URL being shared, it causes an impulse response on the other platforms as well as on itself. This effect increments the probability of a new event related to the original one to happen; the re-posting of the URL.
Figure 7.2 depicts an Hawkes model consisting of two processes, namely, Facebook and Twitter. The first event occurs on Facebook and causes a wave impulse on Twitter and on Facebook itself, incrementing their probability of posting the same URL. Note that the first event was caused by the background rate of Facebook. This means that the URL was not seen on any platform studied, but it was seen elsewhere, for example from the website of news outlet or from other platforms. Then, the same news is shared on Twitter, increasing again the probability of causing new events on the processes. The second event might be occurred for two reasons, i.e., for the Twitter background rate or

for the impulse generated by event one. Generally, for each event we compute the magnitudes of all the active impulses, including the background rate. Then, we assign the probability of being the root cause of the event to the higher impulse. In this case, event two is more probably caused by the background rate of Twitter. The third event occurs soon after on Twitter, again. This time it is mostly influenced by the previous event. Finally, the fourth event occurs later, on Facebook, and it is caused by the Facebook background rate.
According to [44], we refine the definition of the event intensity function of each k-th process belonging to our Hawkes model as:

$$\lambda_{t,k} = \lambda_{0,k} + \sum_{k'=1}^{K} \sum_{t'=1}^{t-1} s_{t',k'} \gamma_{k'->k}[t - t'] \tag{7.2}$$

where:

- $\lambda_{0,k}$ is the background rate of process *k*.
- $s \in N^{T \times K}$ is a matrix that keeps track of how many events occur for process *k* at time *t*.
- $\gamma_{k'->k}[t - t']$ is the impulse function. It describes the influence that events happened on platform k' at time t' have on the rate of process *k* at the current time *t*.
- $\sum_{k'=1}^{K}$ combines the effects of all the processes. $\sum_{t'=1}^{t-1}$ considers all the events occurred before current time *t*.

The impulse response function $\gamma(.)$ is composed by a weight matrix $W_{k->k'}$ and a probability mass function $G_{k->k'}[t - t']$. The weight matrix specifies the strength of the connections between process k and process k'. The probability mass function, instead, describes how the interaction changes over time [44].

7.4 METHODOLOGY

Let's now focus on how we organize our experiments. We aim to examine how Facebook and Twitter influence each other, so, we model the posting of URLs using a Hawkes model with K = 2 point processes. For each process, it is possible to influence all the others, as well as itself, that is why it is said to be fully connected [44]. To fit our model, we select all the posts containing an URL occurred on both Facebook and Twitter. For the *Facebook* dataset, we consider only the URLs contained in Table 5.2, because we need to know the exact time at which the news is shared. Then, to have a baseline of normal Facebook behavior to confront with, we also configure a second experiment. This time, we gather the posts containing an URL shared by both

Table 7.1: Total events with at least one URLs shared by both *Facebook* and Twitter and mean background rate for each platform.

		Facebook	Twitter
Events	Mainstream	488	923
	Alternative	657	1102
	Total	1473	2444
Mean λ_0	Mainstream	0.000195	0.001048
	Alternative	0.001483	0.006708

Facebook general and Twitter. Table 7.1 and Table 7.2 show the events gathered for mainstream and alternative news, on each platform, for respectively the first and second experiments.
Next, we compute the following steps to fit our model:

- For each URL u, we create the matrix $s \in N^{T \times 2}$. Here, T is the number of minutes between the first and last occurrence of u, considering both the platforms. Each row of s represents a platform and contains the count of events happened at that particular time instant. We are dividing the occurrences of the same URL in time bins of 1 minute. In this way, events occurring within the same time bin do not influence each other.

- For each URL, we use the computed s matrix to fit an Hawkes model. We follow the approach described in [21, 22] to infer the parameters of the model from the data. The idea is to use Gibbs sampling [4] to obtain the weight matrix ($W_{K \times K}$), background rates (λ_{0,k_1}, λ_{0,k_2}), and shapes of the impulse response functions between the different processes.

- We set the $\Delta t_{max} = 48$ hours, meaning that an occurrence of an event is influential on another one only if the time difference between them is less or equal to 48 hours. We have also made attempts with other values, as 24 and 168 hours, obtaining similar results.

At this step, we have fitted for each unique URL an Hawkes model. Hence, for each of them we now posses the W matrix, background rate, and impulse response functions. We aim to estimate the total impact that Facebook and Twitter have exerted on each other. Therefore, for each URL occurrence, we now want to compute the influence that previous events have enforced on it. This are the steps we follow to compute the total influence:

- For each URL, we create a matrix $probs \in N^{\#occurrences \times K}$. Each row represents an occurrence of the URL. The rows are sorted by increasing time. Columns, instead, represent the platforms

Table 7.2: Total events with at least one URLs shared by both *Facebook general* and Twitter and mean background rate for each platform.

		Facebook	Twitter
Events	Mainstream	9624	2017
	Alternative	2147	388
	Total	12976	3574
Mean λ_0	Mainstream	0.001234	0.00102
	Alternative	0.000061	0.000019

studied. The value in a cell contains the probability of that occurrence being caused by that platform. The values of each row sum up to 1.

- Originally, we initialize the rows of *probs* entering 1 on the platform on which happened that occurrence. This means that each event was caused only by the background rate of the platform, there is no influence yet.

- For each row of *probs*, we now compute the probability of that occurrence being influenced by previous events. For each row of the matrix, except the first one, we compute the time difference from that occurrence to all the previous one.
If $\Delta t = t - t^{'} \leqslant 48$ hours, then the influence enforced by the events happened at time $t^{'}$ on platform $k^{'}$, towards current event at time t occurred on platform k, is equal to:

$$\mathtt{probs}_t = \mathtt{probs}_t * \lambda_{0,k} + \sum_{t'=1}^{t-1} \mathtt{probs}_{t'} W_{k->k'} \gamma(\Delta t, k^{'}, k) \quad (7.3)$$

The idea is to use the strength of the connection between platform k and $k^{'}$, described by the W matrix, and multiply it by the impulse function. The impulse manage the intensity of the influence, giving stronger power to recent events. The result is obtained by repeating this procedure in a recursive way, summing the effects of all the events, happened in a time range of 48 hours, prior to current time t.

- Finally, for each URL, we sum up the columns values of *probs*. In this way, we obtain the influence enforced by the specific platform on the URL occurrences. If we consider only the occurrences originally happened on Facebook and we sum up the values of the Twitter column, we obtain the score for the Twitter-to-Facebook influence. The opposite has to be done to obtain the score for the Facebook-to-Twitter influence.

7.5 CASE STUDY

Let's now focus our attention on one of the most influenced news to see the pattern with which it was shared.
Looking at the results of our influence analysis, we find out that an URL from *ilblogdellestelle.it*, with headline: "Arriva il #DecretoDignità", has one of the highest influence score Twitter-to-Facebook. This means that its occurrences on Facebook were mostly caused by Twitter. Now, we now try to explain why.
First of all, we see that this URL appears 23 times in our datasets; 19 times on Twitter and 4 on Facebook. If we look at the temporal sequence, on June 14, 2018, the URL was shared 7 times on Twitter between 10.53 am and 11.42 am. Then, it makes its first appearance on Facebook, where at 12.05 it was shared two times. On Twitter there are two posts with the same URL again at 12.20 and at 12.25. Next, the news appears two more times on Facebook at 12.25. Finally, it was shared 10 more times on Twitter; 6 times still in the same day, while the others in the following days. Its last appearance it is exactly one month later its first.
It can be seen that this URL is heavily pushed on Twitter before appearing on Facebook. This behavior in reflected in the influence scores of the occurrences on Facebook. Here, the magnitude of the influence Twitter-to-Facebook dominates the Facebook background rate and Facebook-to-Facebook influence.
Overall, our final scores reflect the fact that it is reasonable to think that the accounts that have shared it on Facebook have previously seen it on Twitter. In fact, the influence Twitter-to-Facebook obtained for this specific URL is 3.9 out of a maximum of 4.
Interestingly, this URL shows also signs of the Facebook echo structure previously studied. Indeed, the 4 occurrences on Facebook happen in pairs, and are shared by two different accounts. The same news shared at the same minute, by 2 of the 23 accounts studied, it is a clear evidence of a connection between the accounts. However, considering how we structured our Hawkes model, occurrences that are that near to each other, i.e., in the same minute, are not captured. That is why we cannot see the influence between them.

7.6 RESULTS

We now show the results obtained, in the first experiment, by our influence study. Considering the total number of events (see Table 7.1), there are 899 unique URLs shared by both Facebook and Twitter. With each of them, we fitted an Hawkes model and computed the influence exerted by each platform.
The first results are visible at Figure 7.3. Here, the percentage values show the quantity of events on the destination platform caused by

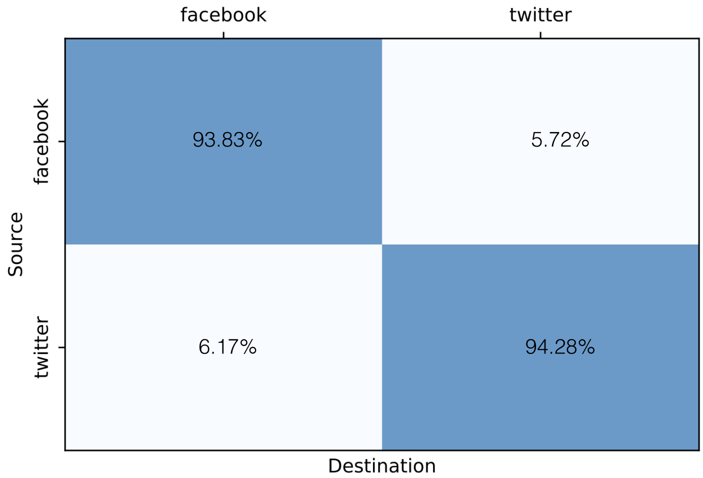

Figure 7.3: First experiment's result. Percentage of destination events caused by the source platform.

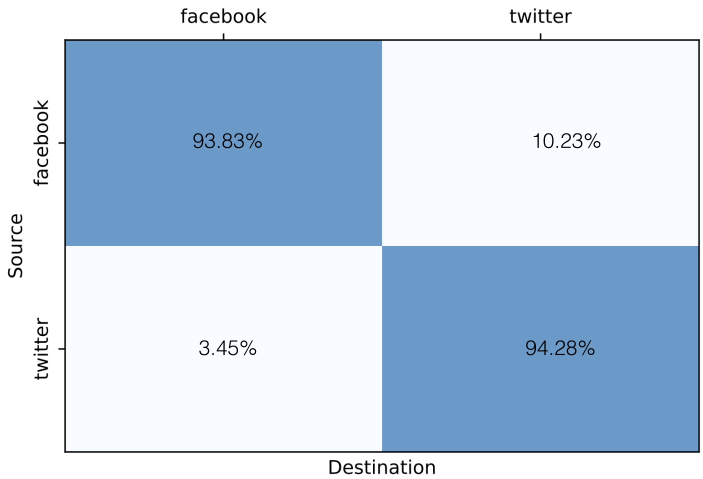

Figure 7.4: First experiment's result. Percentage of destination events caused by the source platform, normalizing by the size of the source platform.

the source platform. As expected, the platforms influence towards themselves is very high. This value capture all the occurrences that happen mostly on only one platform, or those that happen on both, but at distant intervals. Looking at the influence that each platform exerts towards the other, there is equality, with a low prevalence of Twitter over Facebook. We can see that 6.17% of the Facebook posts are actually caused by content previously seen on Twitter, the opposite happens for 5.72% of the cases (tweets).
The second way we report our influence results is visible at Figure 7.4. This time we normalize the influence of each platform by the total number of events happened in the source platform. This results let us get a better insight of the efficiency that each community has, relatively to the number of news they post. While the percentages of the

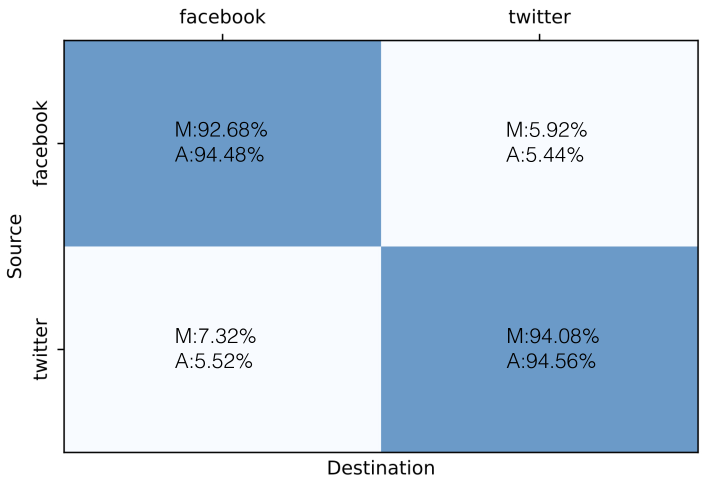

Figure 7.5: First experiment's result. Percentage of destination events caused by the source platform. Comparison between mainstream (M) and alternative news (A).

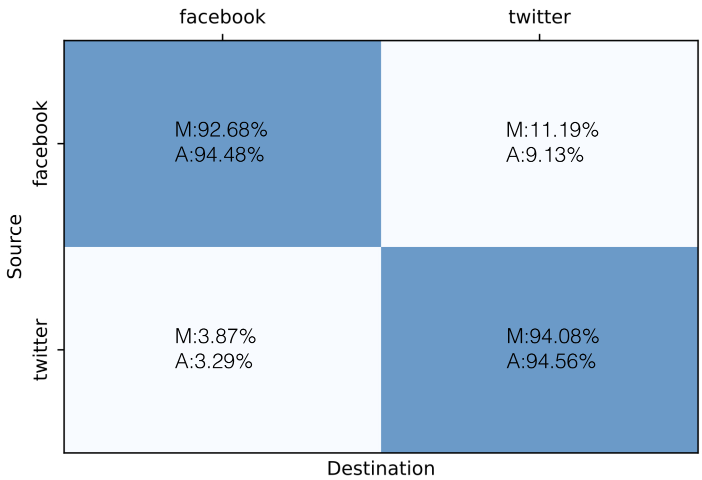

Figure 7.6: First experiment's result. Percentage of destination events caused by the source platform, normalizing by the size of the source platform. Comparison between mainstream (M) and alternative news (A).

influence of the platforms towards themselves have not changed, we can now see an evident difference between the behaviour of Facebook and Twitter. In fact, Facebook is actually more efficient at disseminating news than Twitter. While 10% of Facebook posts cause an event on Twitter, the opposite happens only 3% of the times. This reveal that Twitter is less influential when taking into account the number of news posted.

Next, we want to analyse the difference of traction between the two categories of content; mainstream and alternative news. To do that, we filter the results considering only the the URLs belonging to the specific category we are analysing. From Table 7.1, we notice that the number of events found with an alternative URL is higher for both platforms. This particular behavior is due to the fact that *Facebook* dataset is highly unbalanced towards the alternative side. The high ratio of events to URLs explains the higher background rate for alternative content. Computing again the percentages, we obtain Figure 7.5, where it is shown the comparison between the two categories. The results reveal a difference between alternative and mainstream content. While for alternative news the influence is similar for both the platforms, the number of events caused on Facebook, by Twitter, is higher for mainstream news. Again, it is important to study the efficiency of each platform, considering the size of the source community with respect to the destination one (see Figure 7.6). We can see that, coherently with the previous results, Facebook is more efficient than Twitter for both kinds of content. Considering the number of news posted, Facebook is more influential at pushing mainstream content towards Twitter than alternative, 11% compared to 9%. Twitter instead is equally influential. To asses the statistical significance of this results, we perform a two-sample Kolmogorov-Smirnov test [24] on the influence distributions of mainstream and alternative news. The test confirms the statistical difference ($p<0.01$) between the way alternative and mainstream URLs propagate to Twitter, starting from Facebook. Instead, it proves there is no statistical difference in the way the two types of news spread from Twitter to Facebook ($p>0.01$).

We now perform the second experiment in order to obtain a baseline to compare the previous results with. This time, for the Facebook side, we use the *Facebook general* dataset. Being this dataset originally structured differently from the others, we apply some changes to our model.

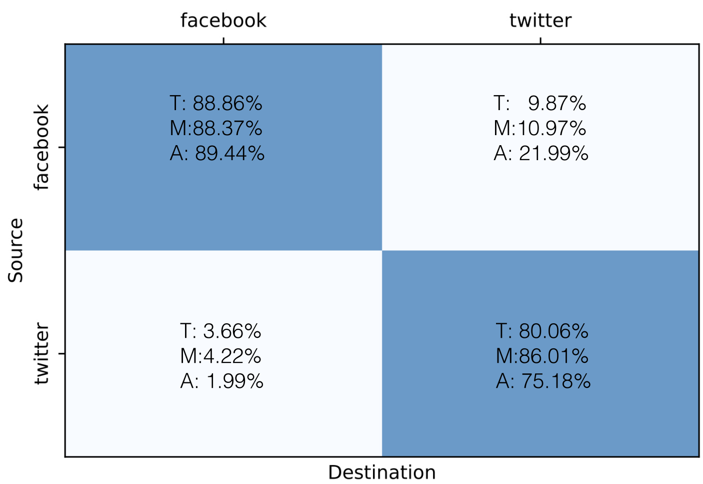

Figure 7.7: Percentage of destination events caused by the source platform comparing mainstream news (M), alternative news (A) and the general behavior (T).

For each URL, we fit an Hawkes model where events occurrences are divided in time bins of 2 hours, instead, than 1 minute. The events gathered for this experiment are resumed at Table 7.2. It is important to notice that, now, the number of mainstream events is higher with respect to alternative ones. Furthermore, the total number of events is almost 10 times more than for the previous experiment. This difference are mainly due to the nature and size of *Facebook general*.
The first results obtained are shown at Figure 7.7. *Facebook general* trigger 9.87% of the tweets in our Twitter dataset, while 3.66% of its posts happened because of Twitter. Hence, while the percentage of events caused by Facebook on Twitter is higher than for the previous experiment, the opposite influence is lower. We have seen that both those values wander around 6% using the *Facebook* dataset (See Figure 7.3).
If we focus on the two categories of content, we can obtain some more specific insight. For Twitter, both alternative and mainstream news are almost equally influential. There is a small prevalence of mainstream content over alternative, 4.22% compared to 1.99%. For Facebook, instead, there is an evident difference between mainstream and alternative news. *Facebook general* alternative news cause 21.99% tweets in our dataset, compared to 10.97% caused by mainstream ones. This difference is not observable in our first experiment, where both the categories are equally influential and cause around 5% of the events on Twitter (See Figure 7.5).
Being *Facebook* and *Facebook general* of different sizes, it is more interesting to compare the efficiency of each dataset rather than the influenced events. For this reason, as for the first experiment, the second way we report the results is by normalizing the values by the total number of events in the source platform (see Figure 7.8). Considering the number of tweets gathered, we can see that Twitter is very efficient at disseminating news towards Facebook. This is evident for mainstream news. In fact, 20.16% of the mainstream URLs that appear on Twitter cause an event on Facebook. In the first experiment this behaviour happens only 7.32% of the times. The percentages for efficiency Facebook-to-Twitter, instead, are a lot smaller than before (see Figure 7.4). Overall, only 2.57% of *Facebook general* posts cause an event on Twitter, while for *Facebook* it happens 10.23% of the times. Interestingly, for *Facebook general* is it possible to see a small prevalence in the influence of alternative news with respect to mainstream ones. On *Facebook* instead, we have seen the exact opposite (see Figure 7.6).

Overall, the results of the two experiments show that the 23 Facebook accounts are more efficient at influencing Twitter than a general Facebook dataset. Moreover, the studied accounts are less influenced by Twitter, especially considering mainstream content. Our theory is that they are organized as a sort of close ecosystem for entering

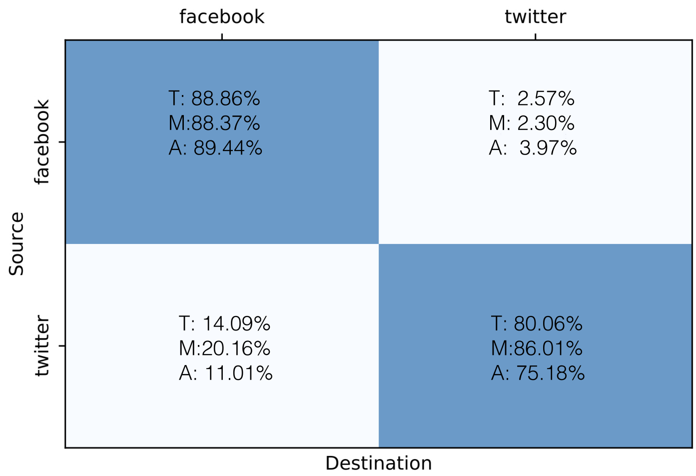

Figure 7.8: Normalized by the size of the source platform (right). We are analysing Facebook general and Twitter.

news. Finally, if we look at the specific type of content spread, we can notice a small prevalence of the efficiency of mainstream content over alternative one. While this behaviour happens in our *Facebook* dataset, the exact opposite is shown in the second experiment using the *Facebook general* dataset.

7.7 TAKE-AWAYS

Summarizing, our influence study brought to light the following results:

1. The studied Facebook accounts exercise an efficient influence at disseminating news towards Twitter. The results show that their efficiency is higher than the one computed on a general set of Facebook behaviour.
2. The Facebook accounts behave as a "close on entry" system. This behaviour is shown by the small number of events caused by Twitter on them.
3. Looking at the specific type of content spread, the Facebook accounts are particular influential at disseminating mainstream news. The exact opposite is shown to be the trend for the general Facebook behaviour.

8 LIMITATIONS AND FUTURE WORK

In this chapter we discuss some of the main limitations we have faced during our study. Finally, we propose further improvements as a guideline for future works.

8.1 LIMITATIONS

In the classification of news domains, we have assigned alternative and mainstream labels to each news outlet. Our goal was to obtain a label for the overall domain behaviour and later use that score to classify all the URLs from that specific domain. This poses a problem: Is the label generalising the content of a single news item?
We had initially verified that some URLs, classified as alternative, actually contained misinformation, after having verified debunked fake news. We then assume that result for all alternative-classified domains. On the other hand, for the mainstream side, we completely trusted the process without employing a further verification.
A second limitation is related to the total number of classified domains. Through our classification process, we managed to classify 38% of all the news domains that appear in our datasets. Considering only those appearing in our *Facebook* dataset, this percentage increases to 54%. In the end, the classified domains let us label as mainstream or alternative 71% of all the unique URLs in our datasets. Then, when we refer to mainstream or alternative news in our studies, we are always referring to that 71% of the total.
Another limitation is related to the usage of the *Facebook general* dataset. We use this dataset in the second experiment of the influence study, to obtain a baseline to compare the other results with. Due to its different structure, the use of this datasets require some assumptions and has some limitations. First, being originally organized in time intervals of two hours, we have to keep this granularity also in the Hawkes model. In fact, for the second experiment, the URL occurrences are divided in time bins of two hours. Second, we re-sample the dataset using a 1% rate. This process allows us to work with the same order of magnitude of Twitter. The sampling rate is chosen for its correspondence with the one used for the Twitter crawler.

8.2 FUTURE WORK

An interesting future development is related to the way news are tracked. In this paper, we focused our attention on the dissemination of news through textual posts. This is a way to analyse the problem, but it is not the only one. In recent years, new types of content have gain traction. Nowadays, is very common to use stories, images, videos, gifts, or memes to communicate. Therefore, most news outlets have adapted to this new situation, being present in more social media platforms than just those based on text content. Instagram for example, is completely based only on images, videos, and stories sharing. It could be very interesting to analyse how this new types of content are changing the overall panorama, and how they affect the news dissemination.

In the end, the topics of disinformation and news dissemination are always actual. Another next possible step for our work could be to integrate our data with new data, regarding other cases. We focused our work on the 2018 Italian elections, so, having available new data about 2020, it could be interesting to see if in a time span of two years the situation is changed.

9 CONCLUSIONS

The research goal was to study the dissemination of alternative and mainstream news during the 2018 Italian elections. We aim to prove the existence of accounts mainly intentioned to spread disinformation. Furthermore, we target the content shared by them, the strategies used in the dissemination of news and their influence towards other social media platforms.

To address the study, we focused our attention on the posts of 23 Facebook accounts. They were originally reported to us by a journalist experienced in fact checking. Then, we created a Twitter dataset made of tweets containing the same keywords and domains included in our Facebook dataset. We classified the news outlets appearing in our datasets into two categories; alternative and mainstream. The classification process uses only publicly available tools and it is mainly based on NewsGuard. We analyse the temporal dynamics of posts. We aim to see how the accounts behavior evolved over time in relation to the two types of content. Next, we moved to the analysis of the text content of posts, keeping a focus on the emotional dynamics and on the hashtags usage. Being, the influence of as many people as possible on the the main goal of malicious users, it is reasonable to presume that the Facebook accounts activity aims to reach other social platforms. Using a statistical approach called Hawkes model, we estimated the influence that their campaigns had on Twitter. Finally, we compared those results with a baseline created ad hoc. Our study leads to the following observations.

The 2018 Italian elections have seen the use of social campaigns to share large quantities of misinformation. In fact, most of the accounts studied are triggered by the approaching of the elections, becoming very active in pushing alternative content just for that specified period. For both Facebook and Twitter, the alternative content is focused in few popular domains. Furthermore, those domains are almost unique for each platform. This means that some of the alternative news domains heavily used to spread misinformation on Facebook aren't, instead, employed on Twitter.

From our temporal analysis, we found out that the implemented campaigns of disinformation are, actually, motivated by real world events. In fact, the quantity of content shared on Facebook increases close to important dates. The election day (04/03/2018) and the day

of sworn for Conte Cabinet (01/06/2018) are the two most important events around which the accounts focused their activity. Moreover, alternative and mainstream news are posted with important differences over time, showing unique patterns. While alternative news are pushed both before and after events, mainstream news are mainly posted after. While alternative news are heavily posted at "sides hours" of the day, i.e., early in the morning or later in the afternoon, mainstream ones, instead, are mostly shared during mid afternoon.
Besides that, on Facebook alternative and mainstream news are consumed differently. Mainstream URLs have a high re-posting behavior in the first minutes after the original post. Alternative URLs, instead, require more time to "be trusted" and re-posted.
In addition, we also found out that 13% of the studied Facebook accounts act as an echo of bigger pages. This behaviour is implemented to increase the visibility of the content originally shared.

Taking a look at the text content of posts, we found out that alternative and mainstream news are associated with different sentiments. This happens for both Facebook and Twitter. We also discovered that generally alternative news are less matched with hashtags. Nevertheless, a deeper analysis on Twitter shows the use of more extremist hashtags for this type of content. For example, the popularity of hashtags such as *#casapound* and *#noiussoli* might represent the intention of share hate-oriented messages.

Finally, our influence studies show that the 23 Facebook accounts are more efficient at disseminating news towards Twitter than a general Facebook dataset. Furthermore, the accounts behave as a "close on entry" system. This means that they are less influenced by the URLs shared on Twitter. Looking at the specific type of content spread, we have seen that the Facebook accounts studied are more efficient at disseminating mainstream content rather than alternative one. This result is in opposition with the baseline, that show that generally Facebook is more efficient at pushing alternative content towards Twitter.

To the best of our knowledge, this work can be considered a first study at the phenomenon of alternative and mainstream news dissemination in the Italian panorama. However, some limitations should be noted. Firstly, our classification process generalise the content of a single news item. Through the pipeline used, we obtain a label for each news outlet. Later, we trust that label for all the URLs belonging to the specific domain. Secondly, we are able to classify 71% of the news URLs gathered in our datasets. 71% is the percentage to which we refer when we target alternative or mainstream content in our analysis. Lastly, the additional Facebook dataset, used during our influence studies, is structured differently from the other datasets. In order to compute the baseline, we have to explicitly define some assumptions which are required to produce comparable results.

Overall, our results show some important insights regarding the phenomenon of the dissemination of disinformation. Our hope is that this results could be used in future works as a starting point for developing some helpful countermeasures. A possible development could be related to the way news are tracked. In this paper, we focused on textual posts. It could be interesting to analyse how new types of media (images, videos, stories) affect the news dissemination.
As the time goes on, it is becoming crucial to find tools to weight the effect of malicious actors on social media and find a way to mitigate their campaigns.

BIBLIOGRAPHY

[1] Avaaz. "Italian Networks Breaking Facebook Rules On Inauthentic Behaviour https://secure.avaaz.org/italynetworks." In: (2019) (cit. on p. 16).

[2] Eytan Bakshy, Itamar Rosenn, Cameron Marlow, and Lada Adamic. "The role of social networks in information diffusion." In: *Proceedings of the 21st international conference on World Wide Web*. 2012, pp. 519–528 (cit. on p. 6).

[3] Samantha Bradshaw and Philip Howard. "The Global Disinformation Order: 2019 Global Inventory of Organised Social Media Manipulation - https://comprop.oii.ox.ac.uk/wp-content/uploads/sites/93/2019/09/CyberTroop-Report19.pdf." In: (2019) (cit. on pp. xv, 2, 5).

[4] George Casella and Edward I George. "Explaining the Gibbs sampler." In: *The American Statistician* 46.3 (1992), pp. 167–174 (cit. on p. 43).

[5] Pew Center. *Election 2016: Campaigns as a direct source of news*. 2016 (cit. on pp. xiv, 2).

[6] Sung Nok Chiu, Dietrich Stoyan, Wilfrid S Kendall, and Joseph Mecke. *Stochastic geometry and its applications*. John Wiley & Sons, 2013 (cit. on pp. 16, 39).

[7] Autorità Per Le Garanzie Nelle Comunicaioni. "News Vs. Fake Nel Sistema Dell'informazione, Interim Report Indagine Conoscitiva Del. 309/16/Cons - https://www.agcom.it/documents/10179/12791486/Pubblicazione+23-11-2018/93869b4f-0a8d-4380-aad2-c10a0e426d83." In: (2018) (cit. on p. 11).

[8] Daryl J Daley and D Vere Jones. *An Introduction to the Theory of Point Processes: Elementary Theory of Point Processes*. Springer, 2003 (cit. on p. 40).

[9] Lizzie Dearden. "Nato accuses Sputnik News of distributing misinformation as part of 'Kremlin propaganda machine' - https://www.independent.co.uk/news/world/europe/sputnik-news-russian-government-owned-controlled-nato-accuses-kremlin-propaganda-machine-a7574721.html." In: (2016) (cit. on p. 13).

[10] Meta Federica. "Elezioni 2018, par condicio anche per Google e Facebook: ecco le linee guida Agcom." In: (2018) (cit. on p. 26).

[11] Emilio Ferrara. "Disinformation and social bot operations in the run up to the 2017 French presidential election." In: *arXiv preprint arXiv:1707.00086* (2017) (cit. on p. xvi).

[12] Pietro Castelli Gattinara, Caterina Froio, and Matteo Albanese. “The appeal of neo-fascism in times of crisis. The experience of CasaPound Italia.” In: *Fascism* 2.2 (2013), pp. 234–258 (cit. on p. 37).

[13] Fabio Giglietto, Laura Iannelli, Luca Rossi, Augusto Valeriani, Nicola Righetti, Francesca Carabini, Giada Marino, Stefano Usai, and Elisabetta Zurovac. “Mapping italian news media political coverage in the lead-up to 2018 general election.” In: *Available at SSRN 3179930* (2018) (cit. on p. 7).

[14] Fabio Giglietto, Nicola Righetti, Giada Marino, and Luca Rossi. “Multi-party media partisanship attention score. Estimating partisan attention of news media sources using Twitter data in the lead-up to 2018 Italian election.” In: *Comunicazione politica* 20.1 (2019), pp. 85–108 (cit. on p. 7).

[15] Georges G Grinstein and Matthew O Ward. “Introduction to data visualization.” In: *Information visualization in data mining and knowledge discovery* 1 (2002), pp. 21–45 (cit. on p. 13).

[16] Jason Horowitz. “Italy, Bracing for Electoral Season of Fake News, Demands Facebook’s Help https://www.nytimes.com/2017/11/24/world/europe/italy-election-fake-news.html.” In: (2017) (cit. on pp. xv, 3).

[17] Cherilyn Ireton and Julie Posetti. *Journalism, fake news & disinformation: handbook for journalism education and training*. UNESCO Publishing, 2018 (cit. on p. 1).

[18] Shih Joseph. “Interpreting The Score And Ratio Of Sentiment Analysis https://www.twinword.com/blog/interpreting-the-score-and-ratio-of-sentiment/.” In: (2018) (cit. on p. 33).

[19] Lee Kevan. “The Biggest Social Media Science Study: What 4.8 Million Tweets Say About the Best Time to Tweet https://buffer.com/resources/best-time-to-tweet-research.” In: (2016) (cit. on p. 29).

[20] Kristina Lerman and Rumi Ghosh. “Information contagion: An empirical study of the spread of news on digg and twitter social networks.” In: *Fourth International AAAI Conference on Weblogs and Social Media*. 2010 (cit. on p. 6).

[21] Scott Linderman and Ryan Adams. “Discovering latent network structure in point process data.” In: *International Conference on Machine Learning*. 2014, pp. 1413–1421 (cit. on p. 43).

[22] Scott W Linderman and Ryan P Adams. “Scalable bayesian inference for excitatory point process networks.” In: *arXiv preprint arXiv:1507.03228* (2015) (cit. on p. 43).

[23] Christopher D Manning, Christopher D Manning, and Hinrich Schütze. *Foundations of statistical natural language processing*. MIT press, 1999 (cit. on p. 14).

[24] Frank J Massey Jr. "The Kolmogorov-Smirnov test for goodness of fit." In: *Journal of the American statistical Association* 46.253 (1951), pp. 68–78 (cit. on pp. 30, 33, 48).

[25] Dipartimento per gli Affari Interni e Territoriali del Ministero dell'interno. "Elezioni trasparenti." In: (2018) (cit. on p. 35).

[26] Robert S Mueller and Man With A. Cat. *Report on the investigation into Russian interference in the 2016 presidential election*. Vol. 1. US Department of Justice Washington, DC, 2019 (cit. on pp. xv, 2).

[27] Alberto Nardelli and Craig Silverman. "Italy's Most Popular Political Party Is Leading Europe In Fake News And Kremlin Propaganda https://www.buzzfeed.com/albertonardelli/italys-most-popular-political-party-is-leading-europe-in-fak." In: (2016) (cit. on pp. xv, 2).

[28] Alberto Nardelli and Craig Silverman. "One Of The Biggest Alternative Media Networks In Italy Is Spreading Anti-Immigrant News And Misinformation On Facebook https://www.buzzfeed.com/albertonardelli/one-of-the-biggest-alternative-media-networks-in-italy-is." In: (2017) (cit. on pp. xv, 2).

[29] "Netpeak https://netpeak.net/about/." In: () (cit. on p. 16).

[30] Bo Pang, Lillian Lee, et al. "Opinion mining and sentiment analysis." In: *Foundations and Trends® in Information Retrieval* 2.1–2 (2008), pp. 1–135 (cit. on p. 33).

[31] Al-Rfou Rami. "Polyglot is a natural language pipeline that supports massive multilingual applications. https://pypi.org/project/polyglot/." In: (2020) (cit. on p. 33).

[32] Alex Reinhart et al. "A review of self-exciting spatio-temporal point processes and their applications." In: *Statistical Science* 33.3 (2018), pp. 299–318 (cit. on p. 40).

[33] Leonard Richardson. "Beautiful soup documentation." In: *April* (2007) (cit. on p. 16).

[34] Marian-Andrei Rizoiu, Young Lee, Swapnil Mishra, and Lexing Xie. "A tutorial on hawkes processes for events in social media." In: *arXiv preprint arXiv:1708.06401* (2017) (cit. on pp. 6, 14, 39–41).

[35] Chengcheng Shao, Giovanni Luca Ciampaglia, Alessandro Flammini, and Filippo Menczer. "Hoaxy: A platform for tracking online misinformation." In: *Proceedings of the 25th international conference companion on world wide web*. 2016, pp. 745–750 (cit. on p. 7).

[36] Elisa Shearer and Jeffrey Gottfried. “News use across social media platforms 2017.” In: *Pew Research Center* 7.9 (2017), p. 2017 (cit. on pp. xiv, 1).

[37] Jacqueline Thomsen. “Mueller: Russia sought to help Trump win but did not collude with campaign.” In: (2019) (cit. on pp. xv, xvi, 2).

[38] Henk C Tijms. *A first course in stochastic models*. John Wiley and sons, 2003 (cit. on p. 40).

[39] Joseph Turian, Lev Ratinov, and Yoshua Bengio. “Word representations: a simple and general method for semi-supervised learning.” In: *Proceedings of the 48th annual meeting of the association for computational linguistics*. Association for Computational Linguistics. 2010, pp. 384–394 (cit. on p. 6).

[40] Michela Del Vicario, Walter Quattrociocchi, Antonio Scala, and Fabiana Zollo. “Polarization and fake news: Early warning of potential misinformation targets.” In: *ACM Transactions on the Web (TWEB)* 13.2 (2019), pp. 1–22 (cit. on p. 11).

[41] Soroush Vosoughi, Deb Roy, and Sinan Aral. “The spread of true and false news online.” In: *Science* 359.6380 (2018), pp. 1146–1151 (cit. on pp. xiv, 1).

[42] Junting Ye and Steven Skiena. “MediaRank: Computational Ranking of Online News Sources.” In: *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2019, pp. 2469–2477 (cit. on p. 11).

[43] Savvas Zannettou, Tristan Caulfield, Jeremy Blackburn, Emiliano De Cristofaro, Michael Sirivianos, Gianluca Stringhini, and Guillermo Suarez-Tangil. “On the origins of memes by means of fringe web communities.” In: *Proceedings of the Internet Measurement Conference 2018*. 2018, pp. 188–202 (cit. on p. 41).

[44] Savvas Zannettou, Tristan Caulfield, Emiliano De Cristofaro, Nicolas Kourtelris, Ilias Leontiadis, Michael Sirivianos, Gianluca Stringhini, and Jeremy Blackburn. “The web centipede: understanding how web communities influence each other through the lens of mainstream and alternative news sources.” In: *Proceedings of the 2017 Internet Measurement Conference*. 2017, pp. 405–417 (cit. on pp. 6, 11, 30, 41, 42).

[45] Savvas Zannettou, Tristan Caulfield, Emiliano De Cristofaro, Michael Sirivianos, Gianluca Stringhini, and Jeremy Blackburn. “Disinformation warfare: Understanding state-sponsored trolls on Twitter and their influence on the web.” In: *Companion Proceedings of The 2019 World Wide Web Conference*. 2019, pp. 218–226 (cit. on p. 34).

[46] Savvas Zannettou, Tristan Caulfield, William Setzer, Michael Sirivianos, Gianluca Stringhini, and Jeremy Blackburn. "Who let the trolls out? towards understanding state-sponsored trolls." In: *Proceedings of the 10th acm conference on web science*. 2019, pp. 353–362 (cit. on p. 6).

[47] Fabiana Zollo and Walter Quattrociocchi. "Misinformation spreading on Facebook." In: *Complex Spreading Phenomena in Social Systems*. Springer, 2018, pp. 177–196 (cit. on pp. 1, 7).