

# Analyzing Polarization of Social Media Users and News Sites during Political Campaigns

Fabrizio Marozzo · Alessandro Bessi

**Abstract** Social media analysis is a fast growing research area aimed at extracting useful information from social networks. Recent years have seen a great interest from academic and business world in using social media to measure public opinion. This paper presents a methodology aimed at discovering the behavior of social network users and how news sites are used during political campaigns characterized by the rivalry of different factions. As a case study, we present an analysis on the constitutional referendum that was held in Italy on 4th December 2016. A first goal of the analysis was to study how Twitter users expressed their voting intentions about the referendum in the weeks before the voting day, so as to understand how the voting trends have evolved before the vote, e.g., if there have been changes in the voting intentions. According to our study, 48% of Twitter users were polarized towards *no*, 25% towards *yes*, and 27% had a *neutral* behavior. A second goal was to understand the effects of news sites on the referendum campaign. The analysis has shown that some news sites had a strong polarization towards *yes* (unita.tv, ilsole24ore.it and linkiesta.it), some others had a *neutral* position (lastampa.it, corriere.it, huffingtonpost.it and repubblica.it) and others were oriented towards *no* (ilfattoquotidiano.it, ilgiornale.it and beppegrillo.it).

**Keywords** Social media analysis · Public opinion · Online information · News sites · Users' polaritazion · Social networks · Political events

## 1 Introduction

In the last years, the production rate of digital data has increased exponentially, with a great contribution from social networks such as Facebook, Twitter, Qzone and Instagram. The large volumes of data generated and gathered by social media platforms can be used to extract valuable information regarding human dynamics and behaviors [4]. Social media analysis is a fast growing research area aimed at extracting useful information from this big amount of data [31]. For example, it is used for the analysis of collective sentiments [30], for understanding the behavior of groups of people [9][10], and to improve the communication between companies and customers [1].

Recently, there has been a great interest from academic and business world for using social media to measure public opinion [2]. Several researchers have used social media data for predicting election results [14], measuring how public opinion changes after important political debates [13] or studying the effects of social media during important recent historical events (e.g., Arab Spring [20]). Other researchers have examined the impact of social

---

Fabrizio Marozzo  
DIMES, University of Calabria  
via P. Bucci, 41C, 87036, Rende, Italy  
E-mail: fmarozzo@dimes.unical.it

Alessandro Bessi  
Information Sciences Institute, University of Southern California  
Marina del Rey, CA, USA  
E-mail: bessi@isi.edu

media spaces on news consumption [18] and on how information spreads through social networks [24].

This paper presents a methodology aimed at discovering the behavior of social network users and how news sites are used during political campaigns characterized by the rivalry of different factions. The methodology is composed of five steps: *i*) definition of the factions and collection of the keywords associated to a political event; *ii*) collection of all the posts generated by social network users containing one or more keywords defined at first step; *iii*) pre-processing of the posts and creation of the input dataset; *iv*) data analysis and mining; and *v*) results visualization. From one hand, the methodology allows to study the users' polarization before a political event, what arguments they used to support their voting intentions, and if such intentions change in the weeks preceding the vote. On the other hand, the methodology allows to study the effects of news sites on a political event, e.g. how many users used information from news sites to support their voting intentions and what news sites can be considered in favor, against or neutral to a given faction.

Unlike works in literature that classify a post manually [17] or with text mining techniques [8, 7, 33], the methodology exploits keywords contained in a post to classify it in favor of a faction. In this way, a post is classified in favor of a faction only if it shows a clear voting indication for a such faction, otherwise we consider the post as neutral. With regard to studying the polarization of news sites, different works in literature use a direct approach that analyzes the contents of articles published by such news sites to understand their political orientation [35, 12]. Our approach instead uses a novel approach that analyzes how users referred to these news sites in their posts for supporting their voting intentions. Other aspects of novelty of the methodology are some analyses we have proposed such as statistical significance of collected data, mobility flows and polarization prediction.

Although the methodology is able to analyze political events characterized by *n*-factions, in this paper we focus on a subset of political events distinguished by the rivalry of only two factions (i.e., *two-faction political events*). This subset includes salient political events, such as *referenda* that see the opposition of two factions (e.g., in favor of *yes* or *no*) or *ballots* (run-off voting) that see the opposition of two candidates competing for the final victory. These events are characterized by some interesting features that makes them interesting to study: *i*) people show a special attention and sensitivity to these events as they are very important for a nation; *ii*) people present a strong polarization in favor of one of the two factions, and this allows separating them in two distinct groups; *iii*) accurate analysis can be done since each user can choose only between two values.

As a case study, we present an analysis on the constitutional referendum that was held in Italy on 4th December 2016. The Italian voters were asked whether they approve a constitutional law that amends the Italian Constitution to reform the composition and powers of the Parliament of Italy. The main supporter of *yes* was the Democratic Party and its leader and Italian Prime Minister Matteo Renzi, whereas in favor of *no* were the main opposition parties and several citizen committees. The referendum saw a high voter turnout (approximately 65% of voters) and a clear victory of *no* (59% of the expressed preferences).

In the weeks before the referendum, we identified a number of keywords (i.e., hashtags) that were used in Twitter to publish *neutral* posts on the referendum, for supporting either *yes* or *no*. We collected 338,592 tweets (1,165,176 if we also consider retweets) that contained those hashtags from 23rd October (5 weeks before the voting day) to 3rd December 2016 (one day before). The number of Twitter users under analysis is 50,717 (139,066 considering also those who published a re-tweet).

A first goal of the analysis was to study how Twitter users expressed their voting intentions about the referendum in the weeks before the voting day, so as to understand how the voting trends have evolved before the vote, e.g., if there have been changes in the voting intentions. According to our study, 48% of Twitter users were polarized towards *no*, 25% towards *yes*, and 27% had a *neutral* behaviour. Regarding the change of opinion in the weeks preceding the vote, the majority of users categorized as supporters of *no* have never changed during the weeks preceding the vote, while a consistent part of the *neutral* users moved towards *no*. A second goal was to understand the effects of news sites on the referendum campaign. The 22% of tweets contained URLs to news related to the referendum. The analysis has shown that some news sites had a strong polarization towards *yes* (unita.tv, ilsole24ore.it and linkiesta.it), some others had a *neutral* position (lastampa.it, corriere.it,

huffingtonpost.it and repubblica.it) and others towards *no* (ilfattoquotidiano.it, ilgiornale.it and beppegrillo.it).

The structure of the paper is as follows. Section 2 describes the methodology proposed in this paper. Section 3 and Section 4 describe respectively how the methodology has been exploited on the Italian constitutional referendum and which results have been achieved. Section 5 discusses related work. Finally, Section 6 concludes the paper.

## 2 Methodology

Given a political event  $\mathcal{P}$  to be analyzed, five are the main steps of the proposed methodology:

1. Definition of the factions  $F$  and collection of the keywords  $K$  associated to  $\mathcal{P}$ ;
2. Collection of all the posts  $P$  generated by social network users containing one or more keywords in  $K$ ;
3. Pre-processing of  $P$  and creation of the input dataset  $D$ ;
4. Data analysis and mining of dataset  $D$ ;
5. Results visualization.

### 2.1 Definition of the factions $F$ and collection of the keywords $K$ associated to $\mathcal{P}$

The political event  $\mathcal{P}$  is characterized by the rivalry of different factions  $F = \{f_1, f_2, \dots, f_n\}$ . Examples of political events and relative factions are: *i*) municipal election, in which a faction supports a mayor candidate; *ii*) political election, in which a faction supports a party; *iii*) presidential election, in which a faction supports a presidential candidate. In this step, we collect the main keywords  $K$  used by social network users to write posts associated to  $\mathcal{P}$ . The keywords can be divided in different subsets, e.g.,  $K = K_{neutral} \cup K_{f_1} \cup \dots \cup K_{f_n}$  as described below:

- The general keywords that can be associated to  $\mathcal{P}$  but cannot be associated to any factions in  $F$  (i.e., are *neutral*) are assigned to  $K_{neutral}$ .
- For each faction  $f_i \in F$ ,  $K_{f_i}$  contains the keywords used to support  $f_i$ .

In this paper we focus on a subset of political events characterized by the rivalry of only two factions  $F = \{f_1, f_2\}$ . Examples of *two-faction events* are: *i*) *referendum*, in which a faction supports a position (e.g., in favor of *yes* or *no*); *ii*) *ballot* (or run-off voting), in which a faction is one of two candidates competing for the final victory. For these events, the keywords are divided in three subsets,  $K = K_{neutral} \cup K_{f_1} \cup K_{f_2}$ , where  $K_{neutral}$  contains the *neutral* keywords,  $K_{f_1}$  and  $K_{f_2}$  the keywords associated respectively to  $f_1$  and  $f_2$ .

### 2.2 Collection of all the posts $P$ generated by social network users containing one or more keywords in $K$

Through the API provided by social networks, we download all the posts containing one or more *keywords* in  $K$ . The posts are not collected in real time, but downloaded a given time after their publication (e.g., 24 hours). In this way, we are able to get some statistics related to the popularity of a post. For example: *i*) number of *shares*, which indicates how many users shared a post with their friends; *ii*) number of *likes*, which indicates how many users found a post useful. Each collected post has at least one key in  $K$ , but may have also other keywords (*co-keywords*) that are useful to understand the arguments used to support the voting intentions.

### 2.3 Pre-processing of $P$ and creation of the input dataset $D$

The goal of this phase is to pre-process the posts in  $P$  to make them ready for the subsequent analysis. Specifically, after pre-processing each post  $p \in P$  is structured as a tuple  $\langle user, text, timestamp, keywords, statistics, URLs, domains, class \rangle$  where

- *user* is the identification of the user who published  $p$ ;
- *text* is the text of the post;
- *timestamp* is the timestamp indicating when  $p$  was published;
- *statistics* contains some statistic data about  $p$ ;
- *keywords* contains the keywords of  $p$ ;
- *URLs* contains all the URLs present in  $p$ ;
- *domains* contains, for each  $u \in \text{URLs}$ , the corresponding domain names;
- *class* is a label that indicates how a post is classified.

The following operations are performed to pre-process the *keywords*, *URLs* and *domains* fields: *i*) all the keywords are transformed to be lowercase and without accented letters (e.g., IOVOTOSI or iovotosí  $\rightarrow$  iovotosi); *ii*) all the short URLs are changed into the corresponding long URLs (for example larep.it  $\rightarrow$  repubblica.it); *iii*) all the alias domains are changed into a single domain (e.g., beppegrillo.it and beppegrillo.com  $\rightarrow$  beppegrillo.it).

The *class* label is computed by analyzing the keywords of a post. A post may be labeled as one of classes  $\{\textit{neutral}, f_1, f_2, \dots, f_n\}$ . Considering the keywords  $\{K_{\textit{neutral}}, K_{f_1}, \dots, K_{f_n}\}$ , Table 1 reports how a post  $p$  is associated to one of the faction  $f_1, \dots, f_n$  or classified as *neutral*. A post is classified as  $f_i$  if it contains at least one keywords in  $K_{f_i}$ , possibly some keywords in  $K_{\textit{neutral}}$ , but no one in other factions  $\{K_{f_1}, \dots, K_{f_{i-1}}, K_{f_{i+1}}, \dots, K_{f_n}\}$ . A post is categorized as *neutral* if has keywords in  $K_{\textit{neutral}}$  and/or keywords in two or many factions  $\{K_{f_1}, \dots, K_{f_n}\}$ . Although there are other approaches in the literature for classifying a post (e.g., manually or with text mining techniques), through our approach, a post is classified in favor of a given faction by analyzing the keywords contained in it, i.e. only if it shows a clear voting indication for that faction, otherwise we consider the post as *neutral*.

In the case of two-faction events, a post may be labeled as one of three classes  $\{\textit{neutral}, f_1, f_2\}$ . Considering the keywords  $\{K_{\textit{neutral}}, K_{f_1}, K_{f_2}\}$ , Table 2 reports how a post  $p$  is associated to one of the two factions  $f_1, f_2$  or classified as *neutral*. A post is classified as  $f_1$  if it contains at least one keywords in  $K_{f_1}$  and possibly some keywords in  $K_{\textit{neutral}}$ . Similarly, a post is classified as  $f_2$  if it contains at least one keywords in  $K_{f_2}$  and possibly some keywords in  $K_{\textit{neutral}}$ . A post is categorized as *neutral* if it has keywords in  $K_{\textit{neutral}}$  and/or keywords in all the two factions  $\{K_{f_1}, K_{f_2}\}$ .

Table 1: Classification of a post by analyzing its keywords in an  $n$ -factions event.

$K_{\textit{neutral}}$	$K_{f_1}$	...	$K_{f_n}$	Class
-	X	-	-	$f_1$
X	X	-	-	$f_1$
-	-	-	X	$f_n$
X	-	-	X	$f_n$
X	-	-	-	neutral
-	X	-	X	neutral
X	X	-	X	neutral

Table 2: Classification of a post by analyzing its keywords in a two-faction event.

$K_{\textit{neutral}}$	$K_{f_1}$	$K_{f_2}$	Class
-	X	-	$f_1$
X	X	-	$f_1$
-	-	X	$f_2$
X	-	X	$f_2$
X	-	-	neutral
-	X	X	neutral
X	X	X	neutral

#### 2.4 Data analysis and mining of dataset $D$

After having built the input dataset  $D$ , it is analyzed through algorithms and techniques for discovering the polarization of social network users and news sites during political campaigns characterized by the rivalry of different factions. In particular, the main goals of this step are as follows.

1. *Analysis of aggregate data.*  $D$  is analyzed to derive statistics about data and to discover the main arguments used by the different factions whose posts are present in  $\mathcal{P}$ .
2. *Statistical significance of collected data* The goal is to assess the significance of  $D$ .
3. *Temporal analysis.* The goal is to analyze how the number of posts supporting the different factions vary along time.
4. *Polarization of users.* Collected data are analyzed to discover how users are polarized towards the different factions.

5. *Mobility flows*. The evolution of users' polarization is studied in the weeks preceding the political event.
6. *Polarization prediction*. The goal is to predict the polarization of users before the political event.
7. *Polarization of news sites*. Collected data are analyzed to discover how news sites are polarized towards the different factions.

## 2.5 Results visualization

Results visualization is performed by the creation of info-graphics aimed at presenting the results in a way that is easy to understand to the general public, without providing complex statistical details that may be hard to understand to the intended audience. The graphic project is grounded on some of the most acknowledged and ever-working principles underpinning a 'good' info-graphic piece. In particular, we follow three main design guidelines: *i*) preferring a visual representation of the quantitative information to the written one; *ii*) minimizing the cognitive efforts necessary to decoding each system of signs; *iii*) structuring the whole proposed elements into graphic hierarchies [11].

Displaying quantitative information by visual means instead of just using numeric symbols - or at least a combination of the two approaches - has been proven extremely useful in providing a kind of sensory evidence to the inherent abstraction of numbers, because this allows everybody to instantly grasp similarities and differences among values. In fact, basic visual metaphors (e.g., the largest is the greatest, the thickest is the highest) enable more natural ways of understanding and relating sets of quantities [32].

## 3 Case of study: Italian constitutional referendum, 2016

We applied the methodology described in the previous section to the constitutional referendum that was held in Italy on 4th December 2016. The Italian voters were asked whether they approve a constitutional law that amends the Italian Constitution to reform the composition and powers of the Parliament of Italy, as well as the division of powers between the State, regions, and administrative entities<sup>1</sup>. The main supporter of the referendum (i.e., in favor of *yes*) was the Democratic Party (in Italian Partito Democratico, PD) and its leader and Italian Prime Minister Matteo Renzi, on the other hand, in favor of *no* the main opposition parties (e.g., Movimento 5 Stelle, Forza Italia) and different citizen committees. The referendum saw a high voter turnout (approximately 65% of voters) and a majority of the votes opposed to the reform (i.e., voting *no*), which exceeded 59% of the expressed preferences. A political effect of the referendum's result was the resignation of the Italian prime minister.

The political event under analysis  $\mathcal{P}$  is a two-faction event  $F = \{yes, no\}$ . We collected the main keywords  $K$  used as hashtags in tweets related to  $\mathcal{P}$ . Such keywords have been grouped as follows:

- $K_{neutral} = \{\#referendumcostituzionale, \#siono, \#riformacostituzionale, \#referendum, \#4dicembre, \#referendum4dicembre\}$
- $K_{yes} = \{\#bastaunsi, \#iovotosi, \#italiachedicesi, \#iodicosi, \#leragionidelsi\}$
- $K_{no} = \{\#iovotono, \#iodicono, \#bastaunno, \#famiglieperilno, \#leragionidelno\}$

Given the keywords  $K = K_{neutral} \cup K_{yes} \cup K_{no}$ , we collected 338,592 tweets containing at least one of these keywords posted from 23rd October (5 weeks before the voting day) to 3rd December 2016 (one day before). The tweets were not collected in real time, but with a delay of 24 hours after their publications so as to capture: *i*) the number of *retweet*, which indicates how many users shared a tweet with their friends; *ii*) the number of *favorites*, which indicates how many users found a tweet useful.

Collected tweets were pre-processed as described in Section 2.3. For instance, Table 3 shows 3 tweets published by a user  $u_i$  before the voting day (translated in English for the Reader's convenience). For each tweet the main fields have been reported in the table.

<sup>1</sup> <http://www.interno.gov.it/italiani-voto-referendum-costituzionale>

Table 3: Examples of tweets on the Italian constitutional referendum.

Text	Timestamp	Keywords	URLs	Class
Why is important to be well informed on #ReferendumCostituzionale	25 Oct 2016 08:00:00	#Referendum Costituzionale	youtube.com/...	neutral
#IoVotoNO: all the reasons to vote against this reform	15 Nov 2016, 09:00:00	#iovotono	ilfattoquotidiano.it/... ilgiornale.it/...	no
Now, wait the results! #referendum4dicembre #iovotoNO #democrazia	3 Dec 2016 10:00:00	#referendum 4dicembre #iovotono #democrazia	-	no

In the first tweet,  $u_i$  expresses the importance of going to vote by using a *neutral* hashtag (#ReferendumCostituzionale) and including a Youtube URL. This tweet is classified as *neutral*. In the second tweet,  $u_i$  shows his/her dissatisfaction with the reform by using a hashtag supporting *no* (#iovotono) and two news sites for motivating his/her voting intention. It is classified as in favor of *no*. The third tweet contains a *neutral* hashtag (#referendum4dicembre), a hashtag supporting *no* (#iovotono) and a co-hashtag (#democrazia). This tweet is classified as in favor of *no*.

## 4 Analysis and results

### 4.1 Analysis of aggregate data

Table 4 reports some statistics about the tweets collected: 338,592 are tweets, 826,584 are retweets and 987,010 are favorites. Filtering the data, we discovered that 43% of tweets contain co-hashtags (e.g., #democrazia, #renzi) and 22% contain URLs. Co-hashtags are useful to understand the arguments used in favor of one or another position. The URLs allows understanding what news site were used by users to support their voting intentions. The number of users under analysis is 50,717 (139,066 considering also the retweets). Figure 1 shows that more than half (54%) of the users published only one tweet on the referendum, 14% two tweets, 7% three, 4% four and 21% five or more tweets.

Table 4: Statistics about collected tweets.

Filter	#Tweets	#Retweets	#Favorites	Total
None	338,592	826,584	987,010	2,152,186
Contains co-hashtags	146,687	449,198	518,088	1,113,973
Contains URL	74,973	139,417	148,888	363,278




Fig. 1: No. of tweets posted by users.

Table 5 reports some statistics about the main hashtags used for collecting tweets, grouped in *yes*, *neutral* and *no*. Next to each hashtag, the number of tweets, retweets and favorites containing such hashtags are reported. The percentage of tweets published with *yes* or *neutral* hashtags are similar (respectively 23% and 24%), and are both half of those in favor of *no* (53%). We also studied how users have used these hashtags to write their tweets: 88% of tweets contain only one or more hashtags of a group (*yes*, *neutral* or *no*), 11% of tweets contain hashtags of two groups (*yes/neutral*, *no/neutral* or *yes/no*), and 1% of tweets contain hashtags of all the groups (*yes/neutral/no*).

Table 5: Main hashtags related to *yes*, *neutral* and *no*.

Hashtag	#Tweets	#Retweets	#Favorites	Total
#bastaunsi	37,268	94,730	133,774	265,773
#iovotosi	38,373	64,419	95,479	198,273
[All hashtags supporting <i>yes</i> ]	76,257	161,306	231,875	469,445
#referendumcostituzionale	36,283	61,940	68,967	167,191
#siono	14,678	28,958	44,460	88,096
#riformacostituzionale	12,233	29,232	30,248	71,715
#referendum	9,727	26,440	27,241	63,409
#4dicembre	7,028	24,715	29,889	61,633
[All <i>neutral</i> hashtags]	81,764	175,123	205,157	462,050
#iovotono	152,638	379,988	430,268	962,895
#iodicono	26,574	107,669	117,233	251,476
[All hashtags supporting <i>no</i> ]	180,562	490,147	549,972	1,220,684

Table 6 shows the main ten co-hashtags used by Twitter users, divided into *yes*, *neutral* and *no*. We note that, in many cases, users who supported *yes* did it by posting tweets reporting Prime Minister’s statements (e.g., #matteorisponde), the opportunity to improve the political system (e.g., #avanti), or information propagated by opponents (e.g., #lebufaledelno). On the other hand, users supporting *no* posted tweets reporting positions from the political opposition (e.g., #m5s, #salvini), willing to leave the constitution as it is (e.g., #costituzione), or hoping to send the prime minister home (e.g., #renziacasa). The *neutral* co-hashtags highlight topics that were treated during the referendum campaign.

Table 6: Main co-hashtags related to *yes*, *neutral* and *no* tweets.

Category	Co-hashtags
<i>yes</i>	#renzi, #sivainpiazza, #matteorisponde, #leopolda7, #avanti, #midiconoche, #m5s, #matteoreni, #pd, #lebufaledelno,
<i>neutral</i>	#agcom, #serracchiani, #renzi, #pd, #mafiacapitale, #mafia, #accozzagliachi, #bufale, #bastapocochevevo, #themancettacandidate
<i>no</i>	#renzi, #salvini, #m5s, #movimentonesti, #trenotour, #costituzione, #nonrubo, #pd, #renziacasa, #deluca,

## 4.2 Statistical significance of collected data

The goal is to assess the statistical significance of the input dataset. Specifically, we studied whether the Twitter users captured in our analysis were actual voters of the referendum, i.e., whether they were Italian citizens able to vote (at least 18 years old).

From the metadata present in the tweets used in our analysis, we extracted aggregate information on the language used to write them and on the location of users who wrote them. Specifically, from the tweet metadata we analyzed the *lang* field<sup>2</sup>, which is a language identifier corresponding to the machine-detected language of the Tweet text (e.g., “en” for English, “it” for Italian, “und” if no language could be detected). In addition, from the user metadata we analyzed the *location* field<sup>3</sup>, which indicates the user-defined location for the accounts profile (e.g., San Francisco, CA).

By analyzing the metadata described above, we can say that:

<sup>2</sup> Twitter API, <https://dev.twitter.com/overview/api/tweets>

<sup>3</sup> Twitter API, <https://dev.twitter.com/overview/api/users>

- All the tweets under analysis have the *lang* field equals to “it” (Italian). The Italian language is mainly used by Italians who reside in Italy (60 million) or abroad (about 4 million). Italian is used as first language<sup>4</sup> only by a small part of Swiss (about 640,000 people), and a very small part of Croats and Slovenes (about 22.000 people).
- 98% of users who have defined the location in their profile live in Italy.

To further show the statistical value of user locations, in Table 7 we compared the number of Twitter users captured in our analysis with the total number of citizens grouped by Italian regions. There is a strong correlation (Pearson coefficient 0.9) between these sets of data. Similar results are obtained by comparing the number of users and the total number of citizen grouped by Italian cities. Also in this case, as shown in Table 8 there is a very strong correlation between these sets of data (Pearson coefficient 0.96).

These statistics give us strong indications about the users analyzed in our case study: it is highly likely that they are voters of the referendum, that is adult Italians citizen. Regarding the last point, statistics show that 96% of Italian Twitter users are adults<sup>5</sup>.

Table 7: Comparison of the number of users and the total number of citizens grouped by region.

Region	N. of users	N. of citizen
Lazio	4,169	5,893,935
Lombardy	4,129	10,014,304
Campania	1,739	5,840,219
Tuscany	1,628	3,743,370
Emilia-Romagna	1,621	4,447,419
Sicily	1,431	5,055,838
Veneto	1,331	4,907,284
Piedmont	1,186	4,394,580
Apulia	1,174	4,066,819
Sardinia	675	1,654,587
Liguria	671	1,565,566
Calabria	565	1,966,819
Friuli-Venezia G.	449	1,218,068
Marches	396	1,539,316
Abruzzo	380	1,322,585
Umbria	310	889,817
Trentino-S. Tyrol	255	1,061,318
Basilicata	202	571,133
Aosta Valley	76	126,732
Molise	73	310,685

Table 8: Comparison of the number of users and the total number of citizens grouped by cities (only 20 of the major Italian cities).

City	N. of users	N. of citizen
Rome	3,499	2,874,529
Milan	2,221	1,353,467
Naples	747	969,456
Turin	548	885,651
Florence	486	382,346
Bologna	452	388,567
Palermo	348	672,398
Genoa	313	582,870
Bari	215	323,503
Catania	198	312,895
Cagliari	188	154,194
Padua	185	209,475
Venice	177	261,496
Verona	172	257,815
Bergamo	168	120,358
Brescia	159	196,205
Modena	127	184,642
Trieste	125	203,974
Udine	123	99,245
Salerno	119	134,857

### 4.3 Temporal Analysis

Figure 2 shows the time series of the number of tweets published during the five weeks preceding the referendum. The tweets in the figure are classified as supporting *yes* (solid blue line), *neutral* (black dashed line), or *no* (solid red line). A fourth time series on all the tweets is represented as a solid black line.

<sup>4</sup> Italian language, [https://it.wikipedia.org/wiki/Lingua\\_italiana](https://it.wikipedia.org/wiki/Lingua_italiana)

<sup>5</sup> Digital in 2017:Italy, <http://www.assocom.org/wp-content/uploads/2017/02/digital-in-2017-italy-we-are-social-and-hootsuite.pdf>




Fig. 2: Time series of tweets published from 23rd October to 3rd December 2016.

All four time series have a similar growing trends (Pearson coefficients of the *yes*, *neutral* and *no* series versus the all series range from 0.87 to 0.97) and show some peaks in the following dates:

- 29th October: It was the day after a major television confrontation between Matteo Renzi in favor of *yes*, and the former PM Ciriaco De Mita in favor of *no*<sup>6</sup>;
- 12th and 23rd November: Debates and discussions in different cities of Italy in favor of *yes* or *no*;
- 2nd December: The last day to make propaganda before the election silence day (3rd December).

We observe that, during the whole observation period, tweets supporting *no* were more than those supporting *yes* or *neutral*. Statistically, every day the number of tweets supporting *neutral* or *yes* are similar, and they both are half of the tweets supporting *no*.

Figure 3 shows the number of tweets aggregated by week day. The interest on referendum increases from Monday to Friday, and then decreases during the weekend.




Fig. 3: No. of tweets per week day.

#### 4.4 Polarization of users

Polarization of a user  $\rho_u \in [-1, 1]$  is defined as

$$\rho_u = 2 \times \frac{|yes_u|}{|yes_u| + |no_u|} - 1,$$

where  $|yes_u|$  and  $|no_u|$  represent, respectively, the number of tweets published by  $u$  classified as *yes* and *no* [5]. A value of  $\rho_u$  close to 1 means that user  $u$  tends to be polarized towards *yes*, while when  $\rho_u$  is close to  $-1$  it means that user  $u$  is polarized towards *no*. In all the analyses of our paper, we focused on users who showed a strong polarization towards

<sup>6</sup> <http://www.ilgiornale.it/news/politica/de-mita-attacca-renzi-tv-io-cambio-partito-tu-amici-1324745.html>

a given faction. For this reason, we chose a high threshold (0.9) to select users with strong polarization in favor of *yes* or *no*. Specifically, we consider users with  $\rho_u > 0.9$  as polarized towards *yes*, users with  $\rho_u < -0.9$  as polarized towards *no*, otherwise *neutral*. Figure 4 shows the probability density function of the users' polarization. We observe a trimodal distribution, indicating that a group of users are polarized towards *yes*, another one has a *neutral* polarization, and another one polarized towards *no*. Specifically, the 48% of users under analysis have a strong polarization towards *no*, 25% towards *yes*, and 27% are *neutral*.




Fig. 4: Probability density function of the users' polarization.

Figure 5 illustrates production patterns of polarized users. In particular, the figure shows the complementary cumulative distribution function (CCDF) of the number of tweets published by users polarized towards *yes* and towards *no*. Both curves point out very similar production patterns between users polarized towards *yes* and users polarized towards *no* [25]. The number of tweets posted by a user does not depend on its polarization: there are a similar number of users who have published at least  $x$  tweets among users polarized both towards *no* and towards *yes*.




Fig. 5: Complementary cumulative distribution function of the number of tweets published by users polarized towards *yes* and towards *no*.

#### 4.5 Mobility flows

Figure 6 represents the evolution of users' polarization in the five weeks preceding the referendum. To study the mobility flows of users, we restricted our analysis on users who have published at least 5 tweets (i.e., 10,436 users). The figure shows how vary the number

of users polarized towards *yes* (blue circles), the number of users polarized towards *no* (red circles), and the number of *neutral* users (gray circles) in the five weeks preceding the referendum. Arrows in the figure show the percentage of users who after one week are polarized as in the previous week and the percentage of *neutral* users who move towards *yes* or *no*. We do not report the moving from *yes* towards *no* (and viceversa) because they are low numbers (less than 3%). Notice that the number of users under analysis increase from week to week (from 1,520 to 10,436), because by collecting new tweets we are able to categorize new users.

We observe that, over the five weeks preceding the referendum, the vast majority of users polarized towards *yes* and *no* tend to maintain their polarization. The biggest changes occur only among users categorized as *neutral*: 10% of neutral users moves towards the *yes* and 20% towards the *no*.

We can conclude that almost all users polarized towards *no* have not changed position during the weeks preceding the vote, and one fifth of the *neutral* users moved towards *no*. Also supporters of the *yes* were very compact, while a lower number of *neutral* users have moved to *yes*.




Fig. 6: Evolution of users' polarization in the five weeks preceding the referendum day: users polarized in favor of *yes* (blue circles), in favor of *no* (red circles), and *neutral* (gray circles).

#### 4.6 Polarization prediction

The goal of this section is to predict the polarization of users before the referendum day. Different machine learning techniques has been studied to evaluate their appropriateness in the considered domain. Among those, some classification algorithms have been tested and the Random Forest (RF) [6] algorithm was selected as it achieved the best performance in terms of accuracy and recall, with limited model building time. Other research works exploited RF for social media analysis due to its high level of accuracy (e.g., see [36], [15], [26]), [23].

Random Forests have been trained for predicting the polarization that a user will have before the voting day, by using information available  $n$  weeks before the referendum, where  $n$  varies from 5 to 1. Specifically, we trained five Random Forest models (one for each value of  $n$ ), each of them trained from this information:

- The input is composed by aggregate information contained in tweets posted by a user at least  $n$  weeks before the referendum. This information is: *i*) number of tweets containing *yes* hashtags, *ii*) number of tweets containing *no* hashtags, *iii*) number of tweets containing *neutral* hashtags, *iv*) total number of tweets, and *v*) number of hashtags used.
- The class is a label that indicates the final polarization of a user (*yes*, *no* or *neutral*) calculated by our methodology using all the information contained in all the tweets posted by a user (i.e., it does not depend on  $n$ ).

To fine-tune the model, we performed a grid search over the parameters' space and we found that the best results are provided by a Random Forest using the entropy criterion and 128 estimators.

Figure 7 shows the classification performance achieved by RF models at different times. In particular, we show micro and macro averaging [34] of the area under the curve (AUC) computed for the model trained with information available at different times. Results are averaged over 10 Monte Carlo cross validation iterations and indicate that the information available 5 weeks before the referendum day provide a classification performance of  $0.849 \pm 0.006$  (micro AUC) and  $0.83 \pm 0.006$  (macro AUC). Such a classification performance increases with the amount of information available, reaching the value of  $0.962 \pm 0.002$  (micro AUC) and  $0.949 \pm 0.001$  (macro AUC) one week before the referendum day.




Fig. 7: User polarization prediction achieved by a Random Forest model using information posted by users from 5 to 1 week before the referendum day.

#### 4.7 Polarization of news sites

Table 9 reports some statistics about tweets containing URLs from the main Italian news sites. Almost 3/4 of such tweets contain URLs from five of the major news sites: beppegrillo.it (36%), ilfattoquotidiano (17%), repubblica.it (12%), huffingtonpost.it (8%) and corriere.it (5%). Since we have registered a greater presence of tweets supporting *no*, the popularity of news sites was been affected if the magazine has written articles close to the positions of *no*.

Table 9: Top 15 news sites used by Twitter users during the referendum campaign.

Hashtag	#Tweets	#Retweets	#Favorites	Total
beppegrillo.it	4,244	12,990	13,575	30,810
ilfattoquotidiano.it	2,027	8,935	7,495	18,457
repubblica.it	1,468	2,537	2,571	6,576
huffingtonpost.it	957	3,150	2,763	6,870
corriere.it	558	1,083	1,235	2,876
unita.tv	509	1,992	2,716	5,218
ilgiornale.it	482	873	764	2,120
ansa.it	269	668	606	1,543
ilsole24ore.it	268	386	349	1,004
formiche.net	216	243	204	664
movimento5stelle.it	206	709	593	1,508
lastampa.it	189	438	526	1,153
possibile.com	173	804	627	1,604
linkiesta.it	173	428	358	959
affaritaliani.it	143	606	541	1,290

Given a news site  $s$ , we compute its polarization as follows

$$\rho_s = 2 \times \frac{|yes_s|}{|yes_s| + |no_s|} - 1,$$

where  $|yes_s|$  and  $|no_s|$  represent, respectively, the number of tweets classified as *yes* and *no* that contain a URL linking to the news site  $s$ . Figure 8 shows the polarization of the main Italian news sites for each category (*yes*, *neutral* and *no*). The figure highlights that some journals had a strong polarization towards *yes* (unita.tv, ilsole24ore.it and linkiesta.it), some

others had a *neutral* position (lastampa.it, corriere.it, huffingtonpost.it and repubblica.it) and others towards *no* (ilfattoquotidiano.it, ilgiornale.it and beppegrillo.it). This result can be explained in two ways: news sites that for editorial choices have supported the campaign of *yes* or *no*, or readers of a certain news site that for political reasons are closer to a certain position.

Figure 9 shows the evolution of the polarization of four representative news sites over the five weeks preceding the referendum day. The figure clearly indicates that the polarization of news sites do not show relevant changes over time.




Fig. 8: Polarization of the main Italian news sites for each category (*yes*, *neutral* and *no*).




Fig. 9: Time series of the polarization of four representative Italian news sites.

## 5 Related work

In recent years, the use of social media for measuring public opinion has become one of the hot topics in social network research [2]. In particular, two are the main areas of research related to this paper: *i*) the use of social media to measure public opinion and predict election results; *ii*) the impact of social media on news consumption and on how information spreads through social networks. For each of these research areas, the main related work has been described.

Murphy et al. [27] examined the potential impact of social media on public opinion research, as an important way for facilitating and/or replacing traditional survey research

methods. The authors highlighted several problems related to this topic, for example: *i*) not every member of the public uses social network platforms; *ii*) incomplete and not accurate information published by social users; *iii*) legal regulations about data collected. OConnor et al. [29] correlated Twitter data with several public opinion time series; Anstead and O’Loughlin [3], by analyzing the 2010 United Kingdom election, suggested the use of social media as a new way to understand public opinion. Others related work attempted to measure the public’s evolving response to stimuli, examining both short term events such as TV political debates [13] or long term events such as economic downturns [16]. An in-depth survey on this topic can be found in [21].

Hermida et al. [18] have examined the impact of social media on news consumption, based on an on-line survey of 1600 Canadians. The study highlights that social networks are a significant source of news for Canadians: two-fifths of users under analysis said that they received news from users who they follow, while a fifth got information from news organizations and individual journalists who they follow. Lerman and Ghosh [24] studied the spread of information on social networks and if their network structure affect how information is disseminated. Specifically, they extracted the active users and track how the interest in news spreads among them. Howard et al. [20] studied the effects of social media during the Arab Spring<sup>7</sup>. By analyzing users posts from different social networks, the authors have reached three main conclusions: *i*) social media played a central role in guiding the political debates during such event; *ii*) a spike in on-line conversations often preceded major events in the real world; and *iii*) social media helped to accelerate the spreading of news and ideas in the world.

To highlight the level of novelty of the methodology we proposed, in the following we review some of the most related research works by discussing differences and similarities with our work. Ceron et al. [8] used a text analysis approach [19] for studying the voting intention of French Internet users in both the 2012 Presidential ballot and the subsequent legislative election. The authors mainly present the results of their analysis by comparing them with official data and predictions made by survey companies. Very few are the implementation details, e.g. is not clear how the statistical value of data was assessed and how tweets and users were classified. Gruzd and Roy. [17] investigated the political polarization of social network users during the Canadian Federal Election, 2011. A sample of tweets posted by 1,492 Twitter users were manually classified based on their self-declared political views and affiliations. The methodology we proposed allows to classify a user automatically by analyzing the posts he/she published - and the keywords he/she used - in the weeks preceding the vote. Nulty et al. [28] surveyed the European landscape of social media using tweets originating from and referring to political actors during the 2014 European Parliament election campaign. With respect to our paper, these authors do not present a methodology but only a hashtag analysis per languages, political parties and candidates. Burnap et al. [7] used Twitter data to forecast the outcome of UK General Election, 2015. The authors applied an automated sentiment analysis tool for classifying tweets. Differently from this methodology, we classified user posts - and consequently users who wrote these posts - by taking advantage of the keywords related to the political event under analysis. Similar considerations can be made for [33], that analyzed about 100.000 political tweets on 2009 German federal election using a text analysis software. Kagan et al. [22] exploited Twitter data for predicting the electoral results of 2013 Pakistan and 2014 Indian elections. The authors studied how the support for a candidate (or opposition to a candidate) was spreading through Internet. In fact, the diffusion model proposed classifies a user by taking into account also the percentage of his/her neighbors (i.e., friends) that have expressed a positive/negative opinion on a candidate/faction. With respect to this work, our methodology evaluates only the content of posts published by a user, but it could be extended considering the opinion of friends of such user. Wagner [35] studied the 2014 Scottish independence referendum for understanding how local newspapers supported the campaign of the referendum. Specifically, the author has analyzed the political position of two local Scottish newspapers (i.e., The Courier and Evening Telegraph), by counting how many stories were neutral, in favor of, and opposed to Scottish independence. With respect to our work, this is a traditional approach that analyzes the textual content of articles published by a news site. We proposed to evaluate the political position of a news site by analyzing how users referred to such news

<sup>7</sup> [https://en.wikipedia.org/wiki/Arab\\_Spring](https://en.wikipedia.org/wiki/Arab_Spring)

site for supporting their voting intentions. Similar considerations can be made for [12], that analyzed the behavior of four leading German on-line newspapers over a timespan of four years.

In summary, this paper presents a methodology aimed at analyzing the polarization of social network users and news sites during political campaigns characterized by the rivalry of two factions (e.g., referenda and ballots). Unlike works in literature that classify a post manually [17] or with text mining techniques [8, 7, 33], our methodology exploits keywords (e.g., hashtags) contained in a post to classify it in favor of a faction. In this way, a post is classified in favor of a faction only if it shows a clear voting indication for a such faction, otherwise we consider the post as neutral. With regard to studying the polarization of news sites, different works in literature use a direct approach that analyzes the contents of articles published by such news sites to understand their political orientation [35, 12]. Our approach instead uses a novel approach that analyzes how users referred to these news sites in their posts for supporting their voting intentions. Other aspects of novelty of the methodology are some analyses we have proposed:

- *Statistical significance of collected data*, to study the statistical significance of data used in our analysis. It gives strong indications about the users and if they are voters of the political event under analysis.
- *Mobility flows*, to analyze the evolution of users' polarization in the weeks preceding the political event. It allows to study if users maintained the same polarization or if they changed their opinion.
- *Polarization prediction*, to predict the polarization of users before the political event. This allows understanding with what precision the polarization of a user can be predict, using information available some weeks before the vote.

The whole methodology and all its analysis have been applied to a real application case such as the Italian constitutional referendum, 2016. We studied the behavior of 50,717 Twitter users by analyzing the 338,592 tweets posted on the referendum by them in the five weeks preceding the vote. The results demonstrate the applicability of our methodology in discovering the behavior of social network users and how news sites are used during political campaigns.

## 6 Conclusions

Social media analysis is an important research area aimed at extracting useful information from the big amount of data gathered from social networks. Recent years have seen a great interest from academic and business world in using social media to measure public opinion.

This paper presents a methodology aimed at analyzing the polarization of social network users and news sites during political campaigns characterized by the rivalry of different factions. From one hand, the methodology allows to study the users' polarization before a political event, what arguments they used to support their voting intentions, and if such intentions change in the weeks preceding the vote. From the other hand, the methodology permits to analyze the effects of news sites on important political events, that is, how many users used information from news sites and what news sites can be considered in favor, against or neutral to a given faction.

The methodology has been validated with an important case study as the Italian constitutional referendum, 2016. According to our study, 48% of Twitter users were polarized towards *no*, 25% towards *yes*, and 27% had a neutral behavior. Regarding the change of opinion in the weeks preceding the vote, the majority of users categorized as supporters of *no* or *yes* have never changed during the weeks preceding the vote, while a consistent part of the *neutral* users moved towards *no* (20%) and towards *yes* (10%). A second goal was to understand the effects of news sites on the referendum campaign. The analysis has shown that some news sites had a strong polarization towards *yes* (*unita.tv*, *ilsole24ore.it* and *linkiesta.it*), some others had a *neutral* position (*lastampa.it*, *corriere.it*, *huffingtonpost.it* and *repubblica.it*) and others were oriented towards *no* (*ilfattoquotidiano.it*, *ilgiornale.it* and *beppegrillo.it*). The polarization of news sites has remained almost unchanged in the weeks preceding the vote.

## References

1. Were all connected: The power of the social media ecosystem. *Business Horizons* **54**(3), 265 – 273 (2011)
2. Anstead, N., O’Loughlin, B.: Social media analysis and public opinion: The 2010 uk general election. *Journal of Computer-Mediated Communication* **20**(2), 204–220 (2015). DOI 10.1111/jcc4.12102. URL <http://dx.doi.org/10.1111/jcc4.12102>
3. Anstead, N., O’Loughlin, B.: Social media analysis and public opinion: The 2010 uk general election. *Journal of Computer-Mediated Communication* **20**(2), 204–220 (2015)
4. Belcastro, L., Marozzo, F., Talia, D., Trunfio, P.: Big data analysis on clouds. In: S. Sakr, A. Zomaya (eds.) *Handbook of Big Data Technologies*, pp. 101–142. Springer (2017). URL [http://dx.doi.org/10.1007/978-3-319-49340-4\\_4](http://dx.doi.org/10.1007/978-3-319-49340-4_4). ISBN: 978-3-319-49339-8
5. Bessi, A., Coletto, M., Davidescu, G.A., Scala, A., Caldarelli, G., Quattrociocchi, W.: Science vs conspiracy: Collective narratives in the age of misinformation. *PloS one* **10**(2), e0118,093 (2015)
6. Breiman, L.: Random forests. *Machine learning* **45**(1), 5–32 (2001)
7. Burnap, P., Gibson, R., Sloan, L., Southern, R., Williams, M.: 140 characters to victory?: Using twitter to predict the uk 2015 general election. *Electoral Studies* **41**, 230–233 (2016)
8. Ceron, A., Curini, L., Iacus, S.M., Porro, G.: Every tweet counts? how sentiment analysis of social media can improve our knowledge of citizens political preferences with an application to italy and france. *New Media & Society* **16**(2), 340–358 (2014)
9. Cesario, E., Congedo, C., Marozzo, F., Riotta, G., Spada, A., Talia, D., Trunfio, P., Turri, C.: Following soccer fans from geotagged tweets at fifa world cup 2014. In: *Proc. of the 2nd IEEE Conference on Spatial Data Mining and Geographical Knowledge Services*, pp. 33–38. Fuzhou, China (2015). ISBN 978-1- 4799-7748-2
10. Cesario, E., Iannazzo, A.R., Marozzo, F., Morello, F., Riotta, G., Spada, A., Talia, D., Trunfio, P.: Analyzing social media data to discover mobility patterns at expo 2015: Methodology and results. In: *The 2016 International Conference on High Performance Computing and Simulation (HPCS 2016)*. Innsbruck, Austria (2016). To appear
11. Cesario, E., Iannazzo, A.R., Marozzo, F., Morello, F., Riotta, G., Spada, A., Talia, D., Trunfio, P.: Analyzing social media data to discover mobility patterns at expo 2015: Methodology and results. In: *The 2016 International Conference on High Performance Computing and Simulation (HPCS 2016)*. Innsbruck, Austria (2016)
12. Dallmann, A., Lemmerich, F., Zoller, D., Hotho, A.: Media bias in german online newspapers. In: *Proceedings of the 26th ACM Conference on Hypertext & Social Media*, pp. 133–137. ACM (2015)
13. Elmer, G.: Live research: Twittering an election debate. *New Media & Society* **15**(1), 18–30 (2013). DOI 10.1177/1461444812457328. URL <http://dx.doi.org/10.1177/1461444812457328>
14. Franch, F.: 2010 uk election prediction with social media. *Journal of Information Technology & Politics* **10**(1), 57–71 (2013). DOI 10.1080/19331681.2012.705080
15. Gokulakrishnan, B., Priyanthan, P., Ragavan, T., Prasath, N., Perera, A.: Opinion mining and sentiment analysis on a twitter data stream. In: *Advances in ICT for emerging regions (ICTer), 2012 International Conference on*, pp. 182–188. IEEE (2012)
16. Gonzalez-Bailon, S., Banchs, R.E., Kaltenbrunner, A.: Emotional reactions and the pulse of public opinion: Measuring the impact of political events on the sentiment of online discussions. *CoRR abs/1009.4019* (2010)
17. Gruzd, A., Roy, J.: Investigating political polarization on twitter: A canadian perspective. *Policy & Internet* **6**(1), 28–45 (2014)
18. Hermida, A., Fletcher, F., Korell, D., Logan, D.: Share, like, recommend. *Journalism Studies* **13**(5-6), 815–824 (2012). DOI 10.1080/1461670X.2012.664430
19. Hopkins, D.J., King, G.: A method of automated nonparametric content analysis for social science. *American Journal of Political Science* **54**(1), 229–247 (2010)
20. Howard, P.N., Duffy, A., Freelon, D., Hussain, M.M., Mari, W., Maziad, M.: Opening closed regimes: What was the role of social media during the arab spring? SSRN (2011). URL <http://dx.doi.org/10.2139/ssrn.2595096>
21. Jungherr, A.: Twitter use in election campaigns: A systematic literature review. *Journal of Information Technology & Politics* **13**(1), 72–91 (2016). DOI 10.1080/19331681.2015.1132401
22. Kagan, V., Stevens, A., Subrahmanian, V.: Using twitter sentiment to forecast the 2013 pakistani election and the 2014 indian election. *IEEE Intelligent Systems* **30**(1), 2–5 (2015)
23. Kwon, S., Cha, M., Jung, K., Chen, W., Wang, Y.: Prominent features of rumor propagation in online social media. In: *Data Mining (ICDM), 2013 IEEE 13th International Conference on*, pp. 1103–1108. IEEE (2013)
24. Lerman, K., Ghosh, R.: Information contagion: an empirical study of the spread of news on digg and twitter social networks. *CoRR abs/1003.2664* (2010)
25. Lievrouw, L., Gillespie, T., Boczkowski, P., Foot, K.: Materiality and media in communication and technology studies: An unfinished project. *Media technologies: Essays on communication, materiality, and society* pp. 21–51 (2014)
26. Monti, C., Rozza, A., Zappella, G., Zignani, M., Arvidsson, A., Colleoni, E.: Modelling political disaffection from twitter data. In: *Proceedings of the Second International Workshop on Issues of Sentiment Discovery and Opinion Mining*, p. 3. ACM (2013)
27. Murphy, J., Link, M.W., Childs, J.H., Tesfaye, C.L., Dean, E., Stern, M., Pasek, J., Cohen, J., Callegaro, M., Harwood, P.: Social media in public opinion research executive summary of the aapor task force on emerging technologies in public opinion research. *Public Opinion Quarterly* **78**(4), 788 (2014). DOI 10.1093/poq/nfu053. URL <http://dx.doi.org/10.1093/poq/nfu053>
28. Nulty, P., Theocharis, Y., Popa, S.A., Parnet, O., Benoit, K.: Social media and political communication in the 2014 elections to the european parliament. *Electoral studies* **44**, 429–444 (2016)

29. O'Connor, B., Balasubramanyan, R., Routledge, B.R., Smith, N.A.: From tweets to polls: Linking text sentiment to public opinion time series. *ICWSM* **11**(122-129), 1–2 (2010)
30. Pang, B., Lee, L.: Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval* **2**(12), 1–135 (2008). DOI 10.1561/1500000011. URL <http://dx.doi.org/10.1561/1500000011>
31. Talia, D., Trunfio, P., Marozzo, F.: *Data Analysis in the Cloud*. Elsevier (2015)
32. Tufte, E.R.: *The Visual Display of Quantitative Information*. Graphics Press, Cheshire, CT, USA (1986)
33. Tumasjan, A., Sprenger, T.O., Sandner, P.G., Welp, I.M.: Predicting elections with twitter: What 140 characters reveal about political sentiment. *ICWSM* **10**(1), 178–185 (2010)
34. Van Asch, V.: Macro-and micro-averaged evaluation measures. Tech. Rep. (2013)
35. Wagner, J.P.: The media and national identity: Local newspapers coverage of scottish independence during the campaign of the 2014 scottish independence referendum. In: *Dealing with the Local* (2017)
36. Zhang, K., Cheng, Y., Xie, Y., Honbo, D., Agrawal, A., Palsetia, D., Lee, K., Liao, W.k., Choudhary, A.: Ses: Sentiment elicitation system for social media data. In: *Data Mining Workshops (ICDMW)*, 2011 IEEE 11th International Conference on, pp. 129–136. IEEE (2011)