Analyzing voter behavior on social media during the 2020 US presidential election campaign

**Loris Belcastro · Francesco Branda ·
Riccardo Cantini · Fabrizio Marozzo ·
Domenico Talia · Paolo Trunfio**

Abstract Every day millions of people use social media platforms by generating a very large amount of opinion-rich data, which can be exploited to extract valuable information about human dynamics and behaviors. In this context, the present manuscript provides a precise view of the 2020 US presidential election by jointly applying topic discovery, opinion mining, and emotion analysis techniques on social media data. In particular, we exploited a clustering-based technique for extracting the main discussion topics and monitoring their weekly impact on social media conversation. Afterwards, we leveraged a neural-based opinion mining technique for determining the political orientation of social media users by analyzing the posts they published. In this way, we were able to determine in the weeks preceding the election day which candidate or party public opinion is most in favor of. We also investigated the temporal dynamics of the online discussions, by studying how users' publishing behavior is related to their political alignment. Finally, we combined sentiment analysis and text mining techniques to discover the relationship between the user polarity and sentiment expressed referring to the different candidates, thus modeling political support of social media users from an emotional viewpoint.

Keywords Social media analysis · Opinion mining · User polarization · Sentiment analysis · Political events

1 Introduction

In recent years, the growing use of social media is generating an amount of information-rich data never seen before. This data, commonly referred as Big Social Data, can be effectively leveraged by a wide range of techniques aimed at modeling the interactions of users on social media, their collective sentiment and behavior, the dynamics of public opinion, and the patterns of information production [34,10,28,9]. All the knowledge extracted through such techniques allows to outline a precise profile of social users, by describing them from a behavioral and psychological viewpoint, and by modeling their perception of events and public decisions.

This manuscript presents an in-depth analysis of the posts published on Twitter during the 2020 US election campaign, aiming at outlining an accurate view of the political event from different points of view. Specifically, several techniques of topic discovery, opinion mining, and emotion analysis were combined in a unified data analysis workflow for investigating: (i) trending topics and their evolution over time, (ii) users' political alignment and publishing behavior, and (iii) users' sentiment and emotional aspects.

Firstly, we extracted the main discussion topics characterizing the 2020 US election campaign by leveraging the unsupervised approach proposed in [8], which relies on the density-based clustering of the latent representation of trending hashtags. Afterwards, in order to achieve a more accurate representation of social media conversation, we studied the

Loris Belcastro · Francesco Branda · Riccardo Cantini · Fabrizio Marozzo · Domenico Talia · Paolo Trunfio
DIMES Department, University of Calabria, Rende, Italy
E-mail: [lbelcastro, fbranda, rcantini, fmarozzo, talia, trunfio]@dimes.unical.it

weekly evolution of the detected topics, which is useful to understand how online discussion evolves over time.

Secondly, we modeled the political alignment of social media users, in order to understand which candidate or party public opinion is most in favor of in the weeks preceding the election day. For this purpose, we exploited IOM-NN (*Iterative Opinion Mining using Neural Networks*), a neural-based opinion mining methodology we previously proposed in [5]. Specifically, a real-time analysis was carried out during the 2020 US presidential election campaign using data gathered from Twitter, correctly determining Joe Biden's lead over Donald Trump before the election day. The achieved results, publicly available through our university web portal[1], represent a remarkable step forward with respect to previous works present in the literature. In fact, to the best of our knowledge, experimental evaluations are carried out after the end of the considered event, while in our case we have given a proof of the real-time effectiveness of IOM-NN, which leads to the possibility of using it for enhancing or even replacing traditional opinion polls. Furthermore, for the sake of completeness, we extended the results of the real-time analysis by focusing on the main swing states, i.e. those states characterized by a high uncertainty about the winning candidate and for this reason by a marked strategic importance. We assessed the statistical significance of the collected data by studying the age, gender and geographical distribution of Twitter users for understanding whether they can be considered voters of the political event. The obtained results confirm the great effectiveness of our approach, which outperformed the average of the latest opinion polls by correctly identifying the leading candidate before the election day in 10 out of 11 swing states. Furthermore, the polarization information achieved by IOM-NN was also leveraged to investigate the temporal dynamics of social media conversation, with the aim of studying how users' publishing behavior is related to their political alignment, and how it reflected the occurrence of external events like debates or rallies.

Thirdly, we analyzed the relationship between the emotional sphere of Twitter users and their political alignment. In particular, we jointly exploited sentiment analysis and text mining techniques for extracting the sentiment of social media users. Then, we combined this information with the polarization achieved by IOM-NN for investigating how a user refers to the candidates while supporting his/her preferred faction, with respect to a broad spectrum of emotions. This step is useful for understanding how the supporters of a particular candidate express their preference on social media. Specifically, they can praise, as in the case of pro-Trump users, their favorite candidate with positive content that shows emotions like joy and confidence. Alternatively, they may be more likely to discredit the opposing candidate, as in the case of pro-Biden users, by producing negative online content characterized by emotions like anger, disgust and sadness.

The rest of the paper is organized as follows. Section 2 reports the most relevant approaches in computational politics and sentiment analysis present in the literature. Section 3 describes the different techniques combined in the proposed analysis workflow. Section 4 describes the experimental evaluation. Finally, Section 5 concludes the paper.

2 Related work

Computational politics is a research area that involves a set of techniques aimed at analyzing users' behavior during a political event of interest, both modeling and influencing their perception and opinion about facts, events and public decisions. With the rapid growth of social media usage, microblogging platforms have become a rich source of valuable information, which can be effectively exploited for investigating the patterns of information diffusion, the interactions between users and their opinion about a specific faction or candidate [5]. According to a recent survey [21], existing literature on computational politics can be categorized into five classes, as discussed in the following.

Community and user modeling. This class of works focuses on modeling the behavior of social media users from both an individual and collective viewpoint. Many works in this category are related to the analysis of homophily, i.e., the connection of groups of users driven by common interests, which leads to the formation of community structures of like-

---

[1] https://www.unical.it/portale/portaltemplates/view/view.cfm?103993 (text in Italian)

minded people [19,4,16]. Other works focus on modeling political affiliation of social users, exploiting community information for predicting the results of a political event [5,12,39].

Information flow. These works investigate how information flows within the network. Most of them analyze the misinformation spread, trying to detect fake news thus limiting its distortion effects on public opinion [25,13,20]. Other works in this category are also aimed at identifying echo chambers, i.e., situations in which the repetition and sharing of information causes the strengthening of an opinion inside a community [17,2,37].

Political discourse. Works in this category model online discussion from different points of view, taking into account demographic aspects, community structure and information diffusion patterns. Many works are aimed at extracting the main topics of discussion through topic modeling [18,41], or identifying political crisis [24]. Opinion mining techniques can be also exploited for identifying the opinion or mood of social media users about those topics, as users' interactions on social media can affect their political engagement [22,3,31].

Election campaigns. Research contributions in this class are aimed at measuring the engagement of the online audience, enabling large scale opinion polls and the management of the political campaign. In fact, social media provide an effective platform for engaging users in political discussion, which is often used by politicians during the political campaigns [43,23]. Moreover, the analysis of political engagement of social users can accurately forecast the final results of the political event under analysis [5,36].

System design. Works in this category propose a full system design of computation politics systems. As an example, Cambre et al. [7] propose a system design that can help to break the echo chamber effect, moderating the online political discussion, while Dade et al. [14] discuss the relationship between political processes, urban environments and situated technologies.

In this work we use opinion mining and sentiment analysis techniques in order to investigate the polarization of US social media users towards the different candidates involved in the 2020 US presidential election. Starting from this, we identify the emotional state (mood) of social users and its relation with their political orientation. Finally, we exploit the results of polarization analysis in order to forecast the final results.

There are several works in literature that rely on text mining and natural language processing algorithms for investigating the opinion of social users and their collective sentiment towards political candidates or parties. Oikonomou et al. [33] used Textblob[2], a Python library for natural language processing, to predict the outcome of US presidential election in three states of interest (i.e., Florida, Ohio and North Carolina). Wong et al. exploited [42] SentiStrength[3], a lexicon-based sentiment analysis tool, for modeling the political behaviors of users by analyzing tweets and retweets. Alashri et al. [1] analyzed Facebook posts about the 2016 US presidential election with CoreNLP[4] [27], one of the most popular tool for natural language processing, to examine the dynamics between candidate posts and comments they received on Facebook and calculate a score for each political candidate for measuring his/her credibility on a given issue. Singh et al. [38] carried out a comparison among four machine and deep learning algorithms (i.e., TextBlob, Naive Bayes, SVM, and BERT [15]) for sentiment analysis. Authors used the 2020 US presidential election as a case study, finding that the use of BERT leads to the best results.

All of the aforementioned techniques are characterized by several issues related to the use of social media data for predicting the outcome of political events, which are *language barrier*, *misclassification*, *data imbalance* and *reliability* [6]. Consequently, in order to achieve a precise estimate of the political polarization of US citizens, we leveraged the IOM-NN technique, specially designed to overcome these issues [5]: (*i*) it is language-independent, as it uses a hashtag-based bag of words representation; (*ii*) it avoids misclassifications using a high threshold on the polarization probability; (*iii*) it uses randomized class balancing algorithms in order to avoid the learning process being biased towards majority classes; (*iv*) it requires a preliminary study of users' representativeness, in order to understand whether they can be considered voters in the political event under analysis (See Section 4.1.1). Moreover, with respect to state-of-art techniques, IOM-NN allows the classification of a much greater number of tweets and users, due to its incremental and iterative nature, which leads to a better quality and robustness of the results.

---

[2] https://textblob.readthedocs.io/en/dev/

[3] http://sentistrength.wlv.ac.uk/

[4] https://stanfordnlp.github.io/CoreNLP/

3 Analysis Workflow

In this work we present an in-depth analysis of the posts published on Twitter during the 2020 US election campaign, with the aim of outlining an accurate representation of this political event from different perspectives, including users' publishing behavior, discussion topics, political alignment and its relationships with the emotional sphere.

For this purpose, several techniques were combined in a unified analysis workflow, represented in Figure 1, composed of the following steps:

- *Collection of posts*: data are gathered from social media by using a set of keywords related to the considered political event.
- *Classification of posts*: the collected posts are classified in favor of a faction according to the detected political support.
- *Polarization of users*: the classified posts are analyzed for determining the polarization of users towards a faction.
- *Topic discovery*: the collected posts are analyzed in order to identify the politically related discussion topics underlying the conversation on social media, modeling their evolution over time.
- *Temporal analysis*: the temporal dynamics of social media conversation are analyzed and combined with the polarity information of classified posts in order to study users' publishing behavior in relation to their political alignment.
- *Emotion analysis*: the polarized posts are exploited for investigating the relationship between the political orientation of users and the different emotions they expressed in referring to the different candidates.

Among the aforementioned steps, the first three jointly constitute the IOM-NN methodology [5], while the fourth follows the approach to topic detection in social data proposed in [8]. IOM-NN is an opinion mining technique aimed at discovering the political polarization of social media users during election campaigns characterized by the competition of political factions. The methodology relies on an iterative and incremental procedure based on feed-forward neural networks, aimed at discovering the political polarization of social media users by analyzing the posts they publish. An open-source implementation of IOM-NN is available on Github[5].
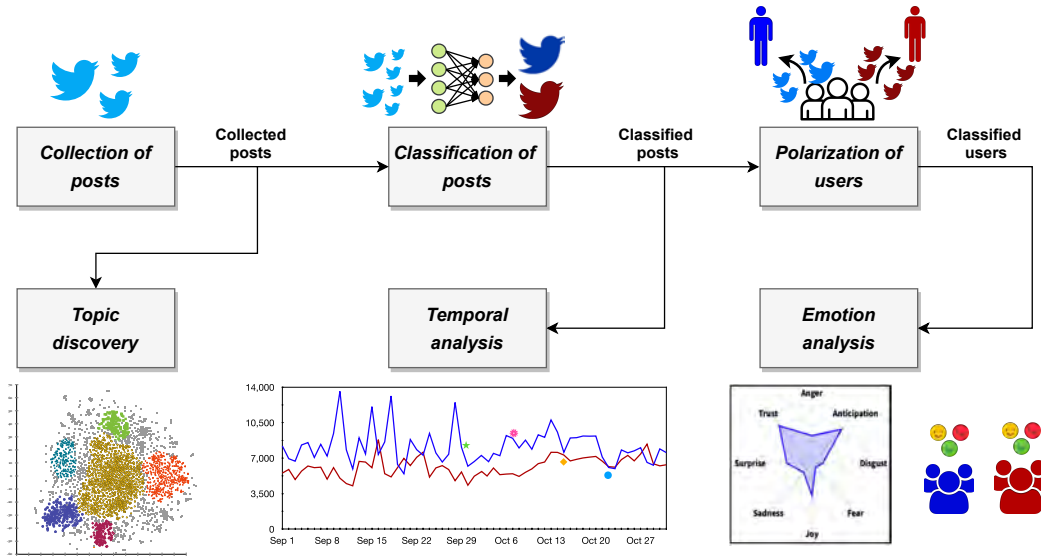

Fig. 1 A graphic representation of our analysis workflow.

In the following sections we provide a detailed description of the proposed analysis workflow. Moreover, to facilitate understanding of the different steps, we will show practical examples by examining a small subset of the collected data.

---

[5] https://github.com/SCAlabUnical/IOM-NN

3.1 Collection of posts

The goal of this step is to collect a set $\mathcal{P}$ of social media posts from different sources (e.g., Twitter), related to the political event $\mathcal{E}$ under analysis. As a first step, the different factions, parties or candidates involved in the political event are identified, defined as the set $\mathcal{F} = \{f_1, f_2, \ldots, f_n\}$. In particular, in the case of the 2020 US presidential election, we focused on the two main candidates Joe Biden and Donald Trump. Afterwards, geotagged posts are gathered by using a set of keywords $\mathcal{K}$ that is partitioned as follows:

- *neutral keywords* ($\mathcal{K}_{context}$) that contains generic keywords that can be associated to $\mathcal{E}$ without referring to any specific faction in $\mathcal{F}$(e.g., #vote, #election2020);
- *faction keywords* ($\mathcal{K}_{\mathcal{F}}^{\oplus} = \mathcal{K}_{f_1}^{\oplus}, \ldots, \mathcal{K}_{f_n}^{\oplus}$) that contains the keywords used for supporting each faction (e.g., #votebiden, #maga).

The keywords selection process requires a small amount of domain knowledge, as these keywords can be manually selected among the trending hashtags that people commonly use to refer to $\mathcal{E}$ on social media. Moreover, the keyword selection process can be automatized by searching for specific patterns, like "#vote + *candidate*", often used for labeling politically polarized posts. We assessed the statistical significance of the collected posts by studying the age, gender and geographical distribution of Twitter users for understanding whether they can be considered voters of the political event. For this purpose we used a wide range of information which can be directly extracted from users metadata (e.g., location and language), or examined starting from statistical reports about the usage of the social media platform in a given country (e.g. user distribution by age and gender). Furthermore, in order to improve the representativeness of the collected posts, user accounts are analyzed, filtering those that show anomalous publishing activity, such as social bots or news sites, or those that have inconsistent information in their profile, such as for example a location that is not defined or does not belong to any of the states considered (see Section 4.1.1).

Collected posts undergo the following preprocessing operations: *i*) the text of each post is converted to lowercase and accented characters are normalized; *ii*) words are lemmatized and stemmed (e.g., vote or votes or voted → vot); *iii*) stopwords are removed; and *iv*) bigrams are identified (e.g., San Francisco → San_Francisco). Figure 2 shows an example of how posts are collected using keywords about the 2020 US presidential election. Some of these keywords are generic (e.g., #vpdebate2020), and others are used to support a specific candidate (e.g., #voteblue for Biden and #trump2020 for Trump).

Fig. 2 Example of how the *collection of posts* step works.

3.2 Classification of posts and temporal analysis

During this step, the posts $\mathcal{P}$ collected in the previous step are classified in favor of a faction by using IOM-NN. Specifically, a preliminary iteration is performed for classifying input posts according to the keywords in $\mathcal{K}_{\mathcal{F}}^{\oplus}$. Posts containing keywords related to exactly one faction are polarized towards that faction, while remaining posts are labeled as neutral. Then, neutral posts undergo an iterative classification process, during which the model exploits the knowledge acquired at the previous iterations. It is worth noticing that, due to the incremental nature of the annotation process, IOM-NN is not tied to a specific set of initial faction keywords and does not require an in-depth knowledge of the political event

under consideration. In fact, even starting from a small but representative set of faction keywords, IOM-NN is able to infer new classification rules iteratively, which implies a good robustness and generalizability of the methodology. Figure 3 shows a classification example of a small set of tweets about the 2020 US presidential election, which exploits the following faction keywords.

- $\mathcal{K}_{Biden}^{\oplus} = \{\#voteblue, \#backtheblue, \#votebiden, \dots\}$;
- $\mathcal{K}_{Trump}^{\oplus} = \{\#votered, \#trump2020, \#maga, \dots\}$.

At iteration 0, IOM-NN uses the keywords in $\mathcal{K}_{\mathcal{F}}^{\oplus}$ for classifying 5 tweets. In the subsequent iterations, the neural model iteratively exploits the tweets classified in the previous steps for generating new classification rules based on co-hashtag relationships. As an example, at iteration 1 the model is trained with the tweet classified at iteration 0, discovering new political-oriented topics of discussion and generating the following classification rules:

- tweets with keywords *#bountygate* are classified in favor of Biden since Donald Trump was accused of paying Moscow's secret agents for the killing of US servicemen in Afghanistan;
- tweets with keywords *#crookedbiden* are classified in favor of Trump since Hunter Biden (i.e., second son of US President Joe Biden), was accused by Donald Trump of wrongdoing in regards to China and Ukraine.

This learning process iterates until the algorithm is no longer able to generate new classification rules and therefore to identify the polarization of new tweets.
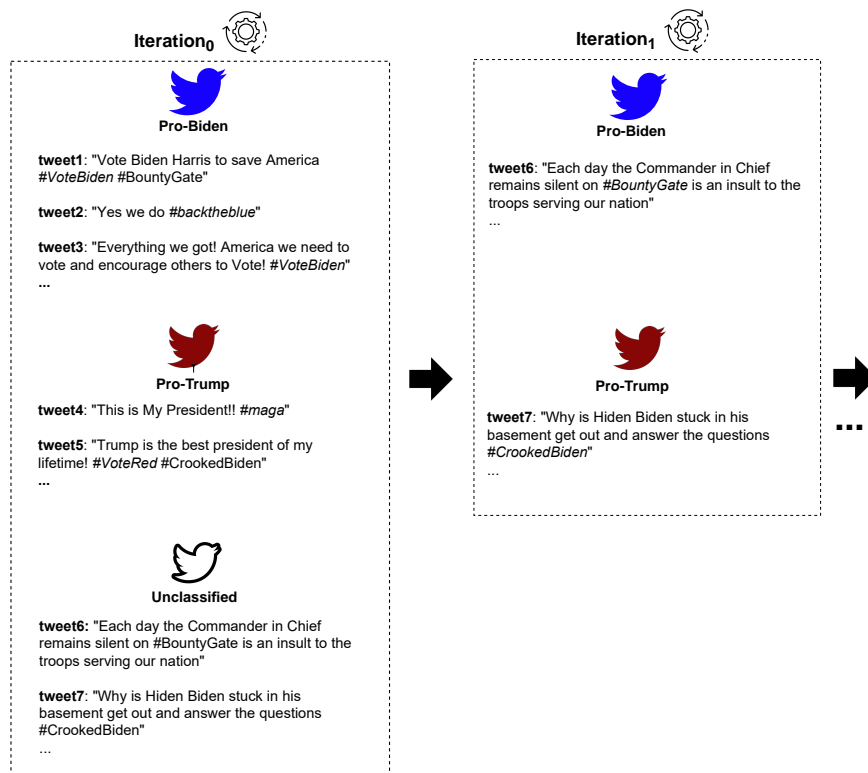

Fig. 3 Example of how the *classification of posts* step works.

After having classified the posts according to the polarity discovered by IOM-NN, this information is used for investigating how the user publishing behavior is related to their political alignment. Specifically, temporal dynamics of social media conversation are analyzed, studying how information is produced by the supporters of both candidates, and how this reflects the occurrence of external events such as debates and rallies.

3.3 Polarization of users

This step is aimed at analyzing the set of previously classified posts in order to determine the *polarization of users* towards a faction. Specifically, the list of classified posts for each user u is computed, filtering out those users that published a number of posts below a given threshold. Afterwards, a score vector v_s^u for each user u is computed, which contains his/her score for each faction. Finally, IOM-NN calculates the overall faction score as the normalized sum of the score vectors. Figure 4 shows how the *polarization of users* step works on the classified posts reported in Figure 3. For each user, the posts in favor of Biden and Trump are counted, discarding those users who have published less than two tweets. Then, the polarization vector for each user is computed containing the percentage of posts published in favor of his/her preferred faction. Lastly, the final score vector is determined, that contains the overall polarization percentages for the two candidates.
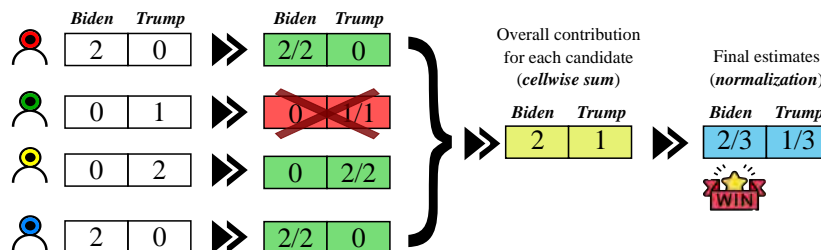

Fig. 4 Example of how the *polarization of users* step works.

3.4 Emotion analysis

This step analyzes the polarized posts for identifying the users' sentiment underlying the online discussion about the presidential candidates. Specifically, we combined the information about the political alignment of social media users with their sentiment and emotional expressions. Firstly, in order to extract the sentiment from online published content, we exploited SentiStrength [40] for the annotation of social media posts. In particular, for each polarized post we computed a positive $\mathcal{S}c^+(p)$ and negative $\mathcal{S}c^-(p)$ sentiment score, both ranging between 1 (neutral) and 5 (strongly positive/negative). Then, the overall sentiment score $\mathcal{S}c(p)$ of a polarized post is obtained as follows: $\mathcal{S}c(p) = \mathcal{S}c^+(p) - \mathcal{S}c^-(p)$. Secondly, we modeled the political orientation of social media users from an emotional point of view by exploiting NRC-EmoLex [30], a publicly available emotion lexicon which has proven its performance in several sentiment and emotion classification tasks, as described in [26], [29], and [32]. Specifically, NRC contains more than 14 thousand English terms labeled by the expressed polarity (i.e., positive or negative) and eight basic emotion categories of Plutchik [35] (i.e., joy, trust, anticipation, sadness, surprise, disgust, fear or anger). Finally, we combined the obtained information with the political alignment discovered by using IOM-NN, in order to extract the overall sentiment and emotions expressed by social media users, while talking about the two candidates.

As an example, Figure 5 and Figure 6 show how the *sentiment analysis* step works on the polarized tweets obtained at the previous step (i.e., a small subset of the collected tweets) for understanding the emotional state of the users who support the different candidates. As we can see, polarized tweets are quite positive for both candidates, but show different emotional profiles. In particular, Biden's supporters show *trust* in the new presidential candidate, while Trump's ones express their *joy* at having Trump as president.

3.5 Topic discovery

This step is aimed at identifying the main politically-related discussion topics characterizing the 2020 US election campaign, by following the unsupervised approach used in [8]. As
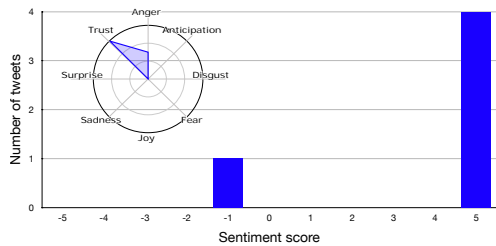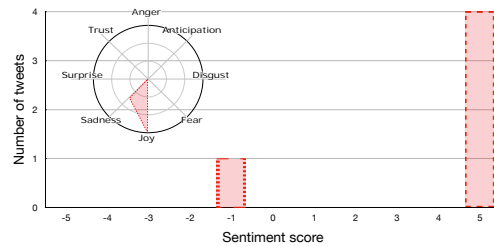
Fig. 5 Example of pro-Biden tweets.

Fig. 6 Example of pro-Trump tweets.

a first step, a Word2Vec model is trained on the entire corpus of tweets, in order to get the latent representation of hashtags and words in a 150-dimensional vector space. We selected the dimension of the embedding space by conducting several experiments, finding out the smallest size for which a clear clustering structure emerged, i.e. the best trade-off between complexity and representativeness. Subsequently, all hashtags are embedded in that 150-dimensional space, whose dimensionality is then reduced by using the t-distributed stochastic neighbor embedding (t-SNE) technique, initialized through principal component analysis (PCA), to obtain a 2D projection of that space. Moreover, in order to reduce noise, all hashtags with a frequency lower than a given threshold are filtered out. Finally, the OPTICS algorithm is used for extracting a clustering structure based on the topic-based separation of hashtags, induced by the projection of their semantic distribution. We have chosen this clustering algorithm due to its ability to discover clusters with arbitrary shape. In addition, compared to classical density-based algorithms such as DBSCAN, it is able to extract clustering structures at different density levels, which in our work is useful for dealing with micro-topics.

4 US 2020 Presidential election analysis and experimental results

In this section we provide an accurate description of the results coming from the analysis of the 2020 US presidential campaign, characterized by a strong rivalry between Joe Biden and Donald Trump. In particular, we analyzed election-related tweets with the aim of outlining a precise representation of this political event from different points of view, in terms of users' publishing behavior, sentiment, political alignment and discussion topics. For this purpose we combined several techniques in an analysis workflow, whose steps are accurately described in Section 3 and whose results are reported in the following sections.

4.1 Data description

The data used to perform the experimental evaluation comes from a public repository that contains a real-time collection of tweets related to the 2020 US presidential election from December 2019 to June 2021 [11]. From such repository we considered only the tweets published close to the election event (from 1 September to 31 October 2020), i.e. about 160 million of which 18 million are tweets (11%), 110 million are retweets (69%), and 32 million are replies (20%), posted by about 29 million users. Only 22% of filtered data contain hashtags (e.g., #trump2020, #bidenharris2020), useful to understand the arguments used in favor of the different candidates. In particular, the percentage of tweets published with at least one hashtag related to Trump (i.e., #trump, #trump2020, and #maga) and Biden (i.e., #bidenharris2020, #biden) is about 31% and 11%, respectively. However, 7% of tweets contain at least one negative hashtag about Trump (i.e., #trumpknew, #pedotrump, #trumphascovid, #trumptaxreturns, #bountygate), whereas only 1% of tweets contain a negative hashtag for Biden (i.e., #crookedjoebiden). In order to ensure the representativeness of the collected posts, we analyzed users' account information, filtering out content posted by users that show an anomalous publishing activity or inconsistent profile information. This step allows to avoid the negative effects caused by the presence of content published by new sites and social bots, which can introduce a heavy bias in social media data [9]. We further analyzed the publishing behavior of the users in the filtered dataset

by determining the Complementary Cumulative Density Function (CCDF) of shared tweets per user.

Specifically, given the random variable X representing the number of shared tweets, it is determined by the frequency of users publishing a number of posts greater than x, i.e. the probability $P(X > x)$.

The scatter plot in log-scale shown in Figure 7, reveals a highly skewed distribution, with few active Twitter users posting a huge amount of tweets, and many users posting infrequently or not at all, the so-called social lurkers.
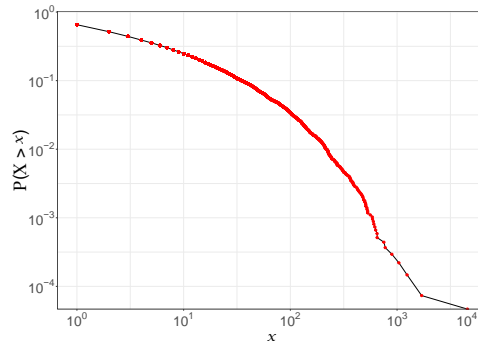

Fig. 7 Complementary Cumulative Density Function (CCDF) of published tweets per user.

4.1.1 Statistical significance of collected data

Here we investigate the statistical significance of the collected data in order to assess users representativeness, i.e., whether they can be considered voters of the political event under analysis.

Firstly, from tweets metadata we extracted aggregate information on the used language of social media users, discovering that most of tweets have the *lang* field set to *English* (about 90%), whereas the remaining 10% is *Undefined* or set to other languages like *Spanish*. Secondly, we compared the number of Twitter users in our dataset, grouped by state, with the number of adult citizens actually living in that state, belonging to the voting-eligible population (VEP)[6]. Specifically, users were associated with states via Twitter metadata, by analyzing the *location* field present in each tweet, which indicates the location defined by the user in his/her Twitter account (e.g., Austin, TX). It is worth noting that, from the textual analysis of this field, it is not always easy to extract a meaningful city/state, as many users either left the field blank, or did not provide precise information (e.g. "USA"), or specified fictitious or non-existent locations (e.g. "the moon" or "NY, Italy"). We measured the strength of this correlation, finding a Pearson coefficient $r = 0.97$, significant at $p < 0.01$. The linear relationship that links users and the voting-eligible population can be easily seen in Figure 8, which depicts an interpolation of the related scatter plot, with a goodness-of-fit $R^2 = 0.93$. Notice that outlier states were not considered in this step in order to achieve meaningful results, by excluding data of different magnitude. In addition, we explored age and gender distribution of analyzed users, finding out that about 94% of them are adults (at least 18 years old)[7] and almost equally divided by gender (22% females and 25% males)[8].

Among all the available tweets we have selected those published by users located in the 11 main swing states (i.e., *Arizona, Florida, Georgia, Michigan, Minnesota, Nevada, New Hampshire, North Carolina, Pennsylvania, Texas, Wisconsin*). We analyzed only these states as they are characterized by a marked political uncertainty and their outcomes have a high probability of being a decisive factor of the electoral event. We made this data, used in all the subsequent analysis steps, publicly available on Github[9].

---

[6] http://www.electproject.org/2020g
[7] https://www.statista.com/statistics/265647/share-of-us-internet-users-who-use-twitter-by-age-group/
[8] https://www.statista.com/statistics/265643/share-of-us-internet-users-who-use-twitter-by-gender/
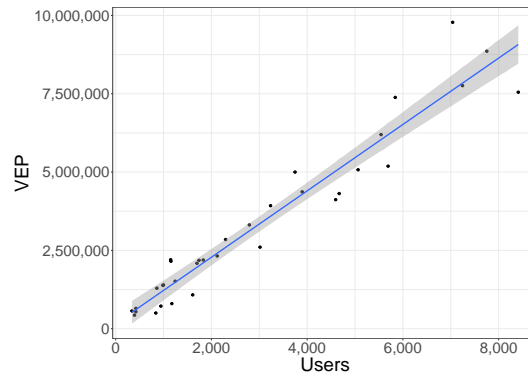[9] https://github.com/SCAlabUnical/USA2020

Fig. 8 Linear interpolation: analyzed users vs. voting-eligible population grouped by US states.

Table 1 reports a comparison between the users we were able to capture for each swing state and the VEP.

Table 1 Number of Twitter users vs voting-eligible population (VEP) grouped by swing states.

State	#Users	#VEP
Arizona	5,692	5,189,000
Florida	16,921	15,551,739
Georgia	5,841	7,383,562
Michigan	8,411	7,550,147
Minnesota	4,596	4,118,462
Nevada	1,156	2,153,915
New Hampshire	1,610	1,079,434
North Carolina	7,245	7,759,051
Pennsylvania	7,040	9,781,976
Texas	19,119	18,784,280
Wisconsin	3,898	4,368,530

The high correlation between the number of analyzed users per state and the VEP leads to a significant set of social media data effectively exploitable to determine the polarization of public opinion However, despite the representativeness of the considered posts, the results achieved by the analysis of the online conversation can be influenced by platform biases. Specifically there exist usage biases due to the distribution of users of a social media platform in terms of gender, age, culture and social status, as well as technical biases related to platform policies about data availability and restrictions imposed in some areas of the world.

4.2 Trending topics of the election campaign

In this step we identified the main politically-related discussion topics characterizing the 2020 US election campaign. Achieved results are shown in Figure 9, where 6 clusters are clearly visible, each one related to a different topic of discussion. Moreover, Table 2 summarizes the discovered topics by reporting the corresponding top hashtags.

The first topic is focused on the criticisms leveled at Trump regarding the management of the health emergency in the United States caused by Covid-19 pandemic. The second one is related to the online discussion about town hall meetings, covering different topics against Trump like discrimination, veterans and climate crisis (e.g., he referred to climate change as a "*hoax*", and to veterans as "*human scum*"). The third one is a general topic about the presidential election. The fourth topic is related to the accusations of corruption and wrongdoing in regards to China and Ukraine leveled against Hunter Biden, i.e., the son of the democratic candidate Joe Biden. The fifth topic focuses on the nomination of the conservative Amy Coney Barrett for a seat on the Supreme Court as successor to the liberal Associate Justice Ruth Bader Ginsburg. Finally, the last topic is related to the online discussion of Trump's supporters, characterized by notorious hashtags like *#maga* or *#kag*.
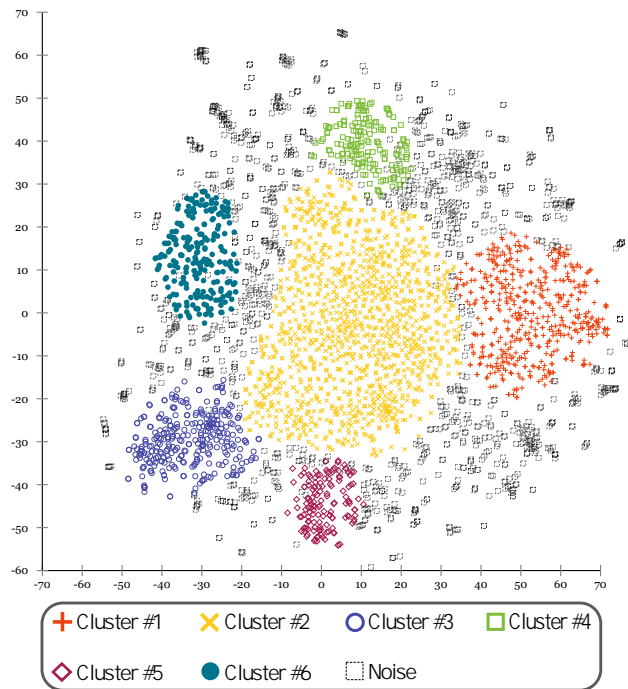
Fig. 9 Unsupervised detection of the main topics underlying the online discussion.

Table 2 Brief description of the identified topics.

Cluster ID	Topic	Top hashtags
#1	Bad management of Covid-19 emergency	#trumpknew, #trumpvirus, #covid, #trumpisaloser, #trumpisanationaldisgrace, #trumpliedpeopledied
#2	Town hall meetings; sub-topics: climate crisis, veterans, discrimination	#cnntownhall, #climatecrisis, #greennewdeal, #respectveterans, #hererightmatters, #stoptrumpsterror
#3	Encouraging people to vote	#election2020, #voteearly, #vote2020, #votebymail, #voteready, #electionday
#4	Accusations against Hunter Biden	#hunterbiden, #bidencrimefamily, #burisma, #ukraine, #hunterbidenemails, #china
#5	US Supreme Court; nomination of Amy Coney Barrett	#scotus, #amyconeybarrett, #filltheseat, #supremecourt, #riprbg, #scotushearings
#6	Support for Trump	#maga, #votetrump2020, #maga2020, #kag, #voteredtosaveamerica2020, #trumppence2020

Once the major discussion topics were detected, we analyzed their overall impact on the online conversation, along with their evolution in the eight weeks included in our observation period, as shown in Figure 10. In particular, we calculated the volume of each hashtag-based topic by determining the percentage of tweets that contain hashtags belonging to the corresponding cluster. Considering our overall observation period, the most relevant topic is about Covid-19 pandemic and it specifically refers to Trump's mismanagement of the health emergency. Other topics are related to the presidential election in general or arise from the publishing activity of Trump's supporters. Also Biden's supporters significantly contributed to the online discussion, by leveraging anti-Trump sub-topics that have emerged from several town hall meetings, about discrimination, veterans and the position of the Republican candidate about the climate crisis.
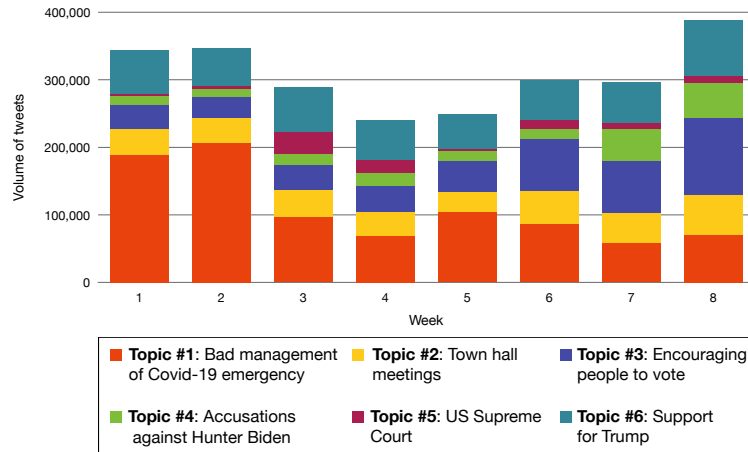
Fig. 10 Weekly volume of tweets related to the detected topics from 1 September to 31 October 2020.

For what concerns the temporal evolution of the detected topics, we found that in the early weeks online conversation focused on the relationship between Trump and Covid-19 pandemic. In addition, the discussion about the US Supreme Court showed a slight increase close to the nomination, announced by Donald Trump, of Judge Amy Coney Barrett as Associate Justice of the US Supreme Court to fill the vacancy left by the death of Ruth Bader Ginsburg. In the following weeks, the focus of the online conversation shifted to various topics related to the approach of the election day and the importance of voting. We also observed an increase in the volume of tweets concerning the accusations leveled against Joe Biden's son (i.e., Hunter Biden), a topic discussed mostly by the Democratic candidate's detractors. Finally, other topics regarding the support voters expressed towards Trump and their criticisms leveled against him linked to town hall meetings showed an almost constant impact on the online conversation.

4.3 Temporal Analysis

In this step we investigated the temporal dynamics of social media conversation, in order to analyze users' publishing behavior, studying how it is related to the detected polarity and how it reflected the occurrence of external events (e.g., debates, rallies, etc.). However, as described by the repository owners in [11], there may be gaps in the dataset due to several issues. Firstly, the data collection step was highly contingent upon the stability of the network and hardware. Secondly, Twitter significantly limits the number of tweets that can be rehydrated. Finally, tweets may no longer be available as users have been removed, banned, or suspended.
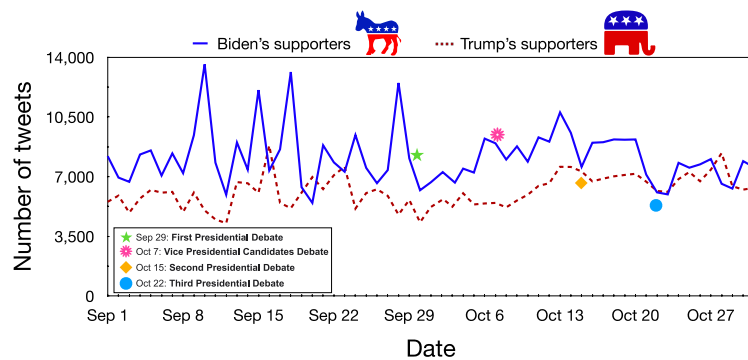

Fig. 11 Time series of polarized tweets published from 1 September to 31 October 2020.

Figure 11 shows the timeline of polarized tweets volume annotated with the four main political debates occurring during the election campaign, i.e., between September 1st and

October 31st. The first observation period (September 1st to September 28th) exhibits significantly different communication dynamics prior to the first debate. Interestingly, this image shows an intense activity spikes of Biden's supporters, as a likely consequence of President Trump's actions:

− *September 10th*: president Trump has attacked Democratic Vice Presidential candidate Kamala Harris.
− *September 15th*: despite being banned by state authorities from holding rallies, President Trump still decided to hold one in Nevada.
− *September 18th*: president Trump blamed blue states for the high number of US Covid-19 fatalities.
− *September 28th*: during a rally in Pennsylvania, Trump called Biden "a dishonest politician and a puppet in the hands of the radical left".

The second and third observation windows (from September 30th to October 31st) show typical weekly cycles of social media chatter, with no particular explosion or shock-related spike from external events, except for October 6th (before the Vice Presidential debate) and October 13th (before the second Presidential debate).

4.4 Comparative analysis with opinion polls

In this step we assessed the effectiveness of our approach in determining the polarity of social media users with the aim of understanding which candidate or party public opinion is most in favor of. A first remarkable result was obtained through a real-time analysis, carried out on Twitter data collected during the two weeks before the election day. Specifically, IOM-NN was able to correctly determine Joe Biden's lead over Donald Trump, especially in Georgia, where a Democratic candidate had not won since 1992 with the election of Bill Clinton. This promising result, publicly available through our university web portal[10], represents a step forward with respect to our previous work, as it gives a clear proof of the real-time effectiveness of IOM-NN, which suggests the possibility of using it to enhance or even replace traditional opinion polls.

Starting from the encouraging real-time results, we extended that analysis by focusing on the main eleven swing states, as described in Section 4.1.1. Specifically, we compared the results obtained through IOM-NN with the average values of the latest opinion polls before the election[11]. For each analyzed state, Table 3 reports the real voting percentages, opinion polls, and IOM-NN estimates. The two candidates (i.e., Joe Biden and Donald Trump) are indicated with "B" and "T", respectively. The winning candidate is written in bold when it is correctly identified.

The results of the comparison is summarized in Figure 12, which shows that the estimates achieved by IOM-NN, related to the voting intentions of social media users are more in line with the actual behaviors of voters with respect to the opinion polls, thus giving a clue to the final result in 10 out of 11 swing states (with an average accuracy of 91%).
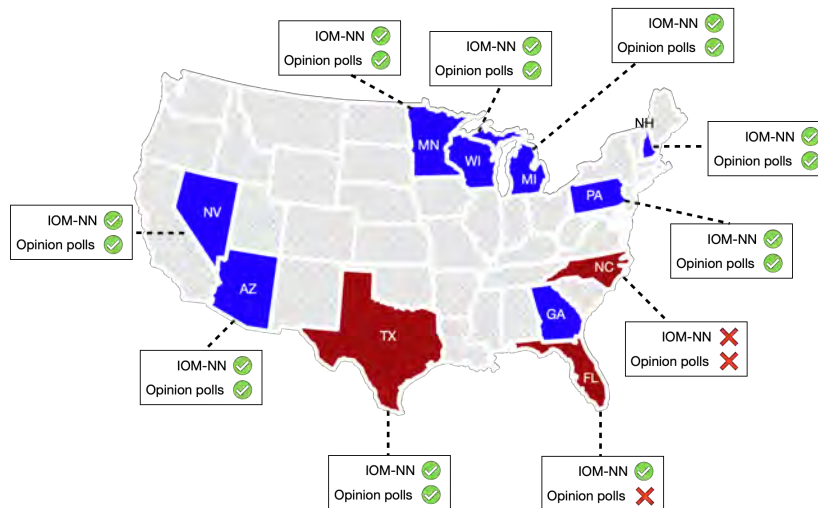
Using this metric we penalize the inversions of polarity which can be a crucial issue while analyzing these kinds of states characterized by a high degree of uncertainty. Notice that, for what concerns North Carolina, neither the estimates achieved by IOM-NN nor the opinion polls were in line with the actual outcome in this state. This is a common situation as the results achieved by the polls and IOM-NN must be understood as an estimate of the polarization of public opinion in the weeks preceding the election day, not always in accordance with the actual behavior of voters. Moreover, a noteworthy advantage of IOM-NN with respect to traditional opinion polls, is the ability to capture the opinion of a larger number of people more quickly and at a lower cost. This makes IOM-NN a valid support to enhance or even replace opinion polls, by providing relevant insights useful to understand the dynamics of the election campaign.

---

[10] `https://www2.unical.it/portale/portaltemplates/view/view.cfm?103993` (text in Italian)
[11] https://www.270towin.com/2020-polls-biden-trump/

Table 3 Comparison between voting percentages estimated by IOM-NN and the latest opinion polls.

State	Real percentages		Opinion polls		IOM-NN	
	B	**T**	**B**	**T**	**B**	**T**
Arizona	49.4	49.1	**48.0**	45.8	**50.2**	48.3
Florida	47.9	51.2	48.7	46.0	48.0	**51.1**
Georgia	49.5	49.2	**47.6**	47.4	**52.7**	46.0
Michigan	50.6	47.8	**49.9**	44.4	**55.4**	43.0
Minnesota	52.4	45.3	**51.6**	41.8	**55.1**	42.6
Nevada	50.1	47.7	**49.4**	44.4	**49.8**	48.0
New Hampshire	52.7	45.4	**53.4**	42.4	**50.9**	47.3
North Carolina	48.6	49.9	47.8	47.5	56.6	41.9
Pennsylvania	50.0	48.8	**49.4**	45.7	**55.7**	43.1
Texas	46.5	52.1	47.5	**48.8**	46.1	**52.5**
Wisconsin	49.4	48.8	**52.0**	42.8	**56.3**	41.9
Correctly classified	-		**9/11**		**10/11**	
Tweets	-		-		670,451	
Users	-		≈ 11,000		57,116	
Avg. Acc	-		0.82		0.91	

Fig. 12 Comparison between IOM-NN and the latest opinion polls in identifying the winning candidate.

4.5 Emotion Analysis

The goal of this last step is to model the political orientation of Twitter users from an emotional point of view. To this purpose, we used the SentiStrength tool (as explained earlier in Section 3.4), for discovering the existing relationships between user polarity and the sentiment expressed in referring to the two presidential candidates. Then, for each polarized tweet we explored the emotion the tweet conveys. Figure 13 and Figure 14 describe the sentiment and the emotional state of the tweets with the relative intensity of the tweets produced by Trump and Biden supporters, respectively.

What appears evident is that, on average, the tweets produced by Trump's supporters are significantly more positive than those produced by Biden's supporters, which devote a significant number of negative tweets to their opponent.

For what concerns the detected emotions, Trump's supporters express *joy* and *confidence* about Trump, while *fear* about Biden's election. Biden's supporters, instead, show *trust* and *anticipation* in having Biden as future president of the United States, with a more marked presence of negative emotions about Trump, like *anger*, *disgust* and *sadness*.

Tables 4 and 5 show various examples of tweets including in the analysis, showing how our approach can model social media conversation from an emotional point of view.
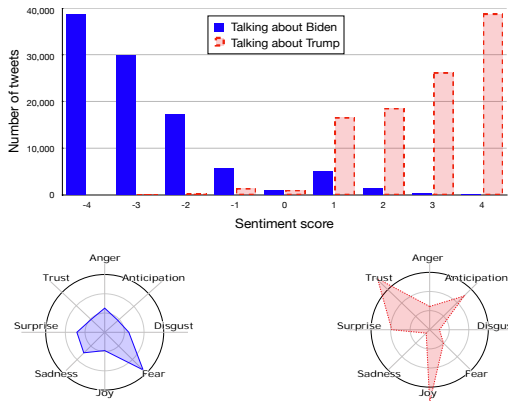
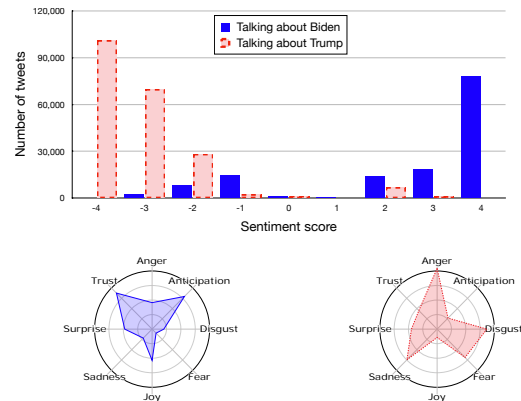**Fig. 13** Distribution of sentiments and emotions of pro-Trump tweets.

Fig. 14 Distribution of sentiments and emotions of pro-Biden tweets.

Table 4 A sample of pro-Trump tweets showing different emotions.

Tweet	About	Sentiment	Emotion
"First time registered voter excited to vote for @realDonaldTrump #FourMoreYears"	Trump	Positive	Joy
"#JoeBiden Democrats support domestic terrorists and exploit race and gender for political gain. I'm afraid for America #PennsylvaniansForTrump"	Biden	Negative	Fear
"You are a disgrace to politicize the death of these people, but obviously you don't care. #JoeBiden #BidenHarris"	Biden	Negative	Disgust
"#realDonaldTrump If anyone can do it, you can. Best President ever! Godspeed sir. #AmericaFirst #MAGA2020"	Trump	Positive	Trust

Table 5 A sample of pro-Biden tweets showing different emotions.

Tweet	About	Sentiment	Emotion
"We all need to #VoteBiden to make this happen #VOTE"	Biden	Positive	Anticipation
"#realDonaldTrump You are a racist and a loser. #TrumpIsALoser #RacistTrump"	Trump	Negative	Disgust
"Today is a sad day. News reports are talking about 200,000 Americans dead from Covid-19 so far #TrumpKnew #COVID19"	Trump	Negative	Sadness
"I want empathy and decency in the White House. #BidenHarris2020ToSaveAmerica"	Biden	Positive	Trust

5 Conclusions and Final Remarks

The widespread use of social media can be exploited to extract useful information concerning people's behaviors and interactions.

In this paper we presented an in-depth analysis of the posts published on Twitter during the 2020 US election campaign, jointly exploiting several techniques for topic discovery, opinion mining and emotion analysis in a unified analysis workflow, with the aim of outlining an accurate representation of this political event from different points of view. In particular, we extracted the main discussion topics following a clustering-based approach, monitoring their weekly impact on social media conversation. Moreover, we leveraged IOM-NN to estimate the polarization of Twitter users regarding the two main candidates Donald Trump and Joe Biden, both in real-time and by focusing on the main US swing states. We also investigated the temporal dynamics of the online discussion, combining it with the polarization information coming from IOM-NN, in order to study how users' publishing behavior reflected external events, over time, in relation to their political orientation. Finally, we ex-

ploited sentiment analysis and text mining techniques to discover the relationship between the user polarization, determined with the aid of IOM-NN, and the sentiment expressed in referring to the different candidates, thus modeling political support of Twitter users from an emotional viewpoint.

Experimental evaluation shows that in the early weeks online conversation focused on the relationship between Trump and Covid-19 pandemic and on the nomination of Judge Amy Coney Barrett as Associate Justice of the US Supreme Court. In the following weeks, instead, the focus of the online conversation shifted to other topics including the accusations leveled to Hunter Biden and the criticism leveled against Trump linked to his position about the climate crisis and veterans. Regarding the political polarization of public opinion, IOM-NN was able to achieve meaningful estimates of the voting intentions of social media users, which makes it a valid solution to go beyond traditional opinion polls, by providing relevant insights useful to understand the dynamics of the election campaign. One major drawback of this approach lies in different possible platform biases, such as usage biases due to the distribution of users of a social media platform in terms of gender, age, culture and social status, as well as technical biases related to platform policies about data availability and restrictions imposed in some areas of the world. Finally, as for the analysis of the emotional state of social users, we found out that the tweets produced by Trump's supporters are significantly more positive than those produced by Biden's supporters. In particular, *i*) Trump's supporters express *joy* and *confidence* about Trump, while *fear* about Biden's election; *ii*) Biden's supporters show *trust* and *anticipation* in having Biden as future president of the United States, with a more marked presence of negative emotions about Trump, like *anger*, *disgust* and *sadness*.

As future work, we will apply the presented analysis workflow to other scenarios, such as product adoption analysis and reputation evaluation of companies. In fact, it can be easily generalized to different use cases, as it is not tied to any specific application domain, and only relies on the representativeness of the analyzed posts. Moreover, we can integrate other techniques in our workflow, introducing new steps aimed at improving the quality of the achieved results. As an example, a hashtag recommendation model can be used for enriching the information content of the analyzed data, since keyword-based approaches like IOM-NN are strongly dependent on the availability of consistent hashtags in social media posts [8].

Acknowledgment

Data Availability Statement

The data that support the findings of this study are publicly available. In particular, this data was gathered using Twitter APIs[12] and is hosted on Github[13].

References

1. Alashri, S., Kandala, S.S., Bajaj, V., Ravi, R., Smith, K.L., Desouza, K.C.: An analysis of sentiments on facebook during the 2016 us presidential election. In: 2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM), pp. 795–802. IEEE (2016)
2. An, J., Quercia, D., Crowcroft, J.: Fragmented social media: a look into selective exposure to political news. In: Proceedings of the 22nd international conference on world wide web, pp. 51–52 (2013)
3. Azarbonyad, H., Dehghani, M., Beelen, K., Arkut, A., Marx, M., Kamps, J.: Words are malleable: Computing semantic shifts in political and media discourse. In: Proceedings of the 2017 ACM on Conference on Information and Knowledge Management, pp. 1509–1518 (2017)

---

[12] `https://developer.twitter.com`

[13] `https://github.com/SCAlabUnical/USA2020`

4. Bastos, M.T., Puschmann, C., Travitzki, R.: Tweeting across hashtags: overlapping users and the importance of language, topics, and politics. In: Proceedings of the 24th ACM conference on hypertext and social media, pp. 164–168 (2013)
5. Belcastro, L., Cantini, R., Marozzo, F., Talia, D., Trunfio, P.: Learning political polarization on social media using neural networks. IEEE Access **8**, 47177–47187 (2020)
6. Bilal, M., Gani, A., Marjani, M., Malik, N.: Predicting elections: Social media data and techniques. In: 2019 international conference on engineering and emerging technologies (ICEET), pp. 1–6. IEEE (2019)
7. Cambre, J., Klemmer, S.R., Kulkarni, C.: Escaping the echo chamber: ideologically and geographically diverse discussions about politics. In: Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems, pp. 2423–2428 (2017)
8. Cantini, R., Marozzo, F., Bruno, G., Trunfio, P.: Learning sentence-to-hashtags semantic mapping for hashtag recommendation on microblogs. ACM Transactions on Knowledge Discovery from Data (TKDD) **16**(2), 1–26 (2021)
9. Cantini, R., Marozzo, F., Talia, D., Trunfio, P.: Analyzing political polarization on social media by deleting bot spamming. Big Data and Cognitive Computing **6**(1) (2022). DOI 10.3390/bdcc6010003
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: 2016 International Conference on High Performance Computing & Simulation (HPCS), pp. 230–237. IEEE (2016)
11. Chen, E., Deb, A., Ferrara, E.: # election2020: the first public twitter dataset on the 2020 us presidential election. Journal of Computational Social Science pp. 1–18 (2021)
12. Chiu, S.I., Hsu, K.W.: Predicting political tendency of posts on facebook. In: Proceedings of the 2018 7th International Conference on Software and Computer Applications, pp. 110–114 (2018)
13. Ciampaglia, G.L., Shiralkar, P., Rocha, L.M., Bollen, J., Menczer, F., Flammini, A.: Computational fact checking from knowledge networks. PloS one **10**(6), e0128193 (2015)
14. Dade-Robertson, M., Taylor, N., Marshall, J., Olivier, P.: The political sensorium. In: Proceedings of the 4th Media Architecture Biennale Conference: Participation, pp. 47–50 (2012)
15. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 (2018)
16. Fraisier, O., Cabanac, G., Pitarch, Y., Besançon, R., Boughanem, M.: Uncovering like-minded political communities on twitter. In: Proceedings of the ACM SIGIR International Conference on Theory of Information Retrieval, pp. 261–264 (2017)
17. Garimella, K., De Francisci Morales, G., Gionis, A., Mathioudakis, M.: Political discourse on social media: Echo chambers, gatekeepers, and the price of bipartisanship. In: Proceedings of the 2018 World Wide Web Conference, pp. 913–922 (2018)
18. Greene, D., Cross, J.P.: Unveiling the political agenda of the european parliament plenary: A topical analysis. In: Proceedings of the ACM web science conference, pp. 1–10 (2015)
19. Grevet, C., Terveen, L.G., Gilbert, E.: Managing political differences in social media. In: Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing, pp. 1400–1408 (2014)
20. Gyongyi, Z., Garcia-Molina, H., Pedersen, J.: Combating web spam with trustrank. In: Proceedings of the 30th international conference on very large data bases (VLDB) (2004)
21. Haq, E.U., Braud, T., Kwon, Y.D., Hui, P.: A survey on computational politics. IEEE Access **8**, 197379–197406 (2020)
22. Hoffmann, C.P., Lutz, C.: Spiral of silence 2.0: Political self-censorship among young facebook users. In: Proceedings of the 8th international conference on social media & society, pp. 1–12 (2017)
23. Hong, S., Nadler, D.: Social media and political voices of organized interest groups: a descriptive analysis. In: Proceedings of the 16th Annual International Conference on Digital Government Research, pp. 210–216 (2015)
24. Keneshloo, Y., Cadena, J., Korkmaz, G., Ramakrishnan, N.: Detecting and forecasting domestic political crises: A graph-based approach. In: Proceedings of the 2014 ACM conference on Web science, pp. 192–196 (2014)
25. Kim, J., Tabibian, B., Oh, A., Schölkopf, B., Gomez-Rodriguez, M.: Leveraging the crowd to detect and reduce the spread of fake news and misinformation. In: Proceedings of the eleventh ACM international conference on web search and data mining, pp. 324–332 (2018)
26. Kiritchenko, S., Zhu, X., Mohammad, S.M.: Sentiment analysis of short informal texts. Journal of Artificial Intelligence Research **50**, 723–762 (2014)
27. Manning, C.D., Surdeanu, M., Bauer, J., Finkel, J.R., Bethard, S., McClosky, D.: The stanford corenlp natural language processing toolkit. In: Proceedings of 52nd annual meeting of the association for computational linguistics: system demonstrations, pp. 55–60 (2014)
28. Marozzo, F., Bessi, A.: Analyzing polarization of social media users and news sites during political campaigns. Social Network Analysis and Mining **8**(1), 1–13 (2018)
29. Mohammad, S.: Portable features for classifying emotional text. In: Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, pp. 587–591 (2012)
30. Mohammad, S.M., Turney, P.D.: Crowdsourcing a word–emotion association lexicon. Computational intelligence **29**(3), 436–465 (2013)
31. 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, pp. 1–9 (2013)
32. Nakov, P., Rosenthal, S., Kiritchenko, S., Mohammad, S.M., Kozareva, Z., Ritter, A., Stoyanov, V., Zhu, X.: Developing a successful semeval task in sentiment analysis of twitter and other social media texts. Language Resources and Evaluation **50**(1), 35–65 (2016)

33. Oikonomou, L., Tjortjis, C.: A method for predicting the winner of the usa presidential elections using data extracted from twitter. In: 2018 South-Eastern European Design Automation, Computer Engineering, Computer Networks and Society Media Conference (SEEDA_CECNSM), pp. 1–8. IEEE (2018)
34. Pang, B., Lee, L.: Opinion mining and sentiment analysis. foundations and trends (r) in information retrieval, 2 (1-2), 1-135 (2008)
35. Plutchik, R.: The nature of emotions: Human emotions have deep evolutionary roots, a fact that may explain their complexity and provide tools for clinical practice. American scientist **89**(4), 344–350 (2001)
36. Saleiro, P., Gomes, L., Soares, C.: Sentiment aggregate functions for political opinion polling using microblog streams. In: Proceedings of the Ninth International C* Conference on Computer Science & Software Engineering, pp. 44–50 (2016)
37. Shu, K., Bernard, H.R., Liu, H.: Studying fake news via network analysis: detection and mitigation. In: Emerging Research Challenges and Opportunities in Computational Social Network Analysis and Mining, pp. 43–65. Springer (2019)
38. Singh, A., kumar, A., Dua, N., Mishra, V.K., Singh, D., Agrawal, A.: Predicting elections results using social media activity a case study: Usa presidential election 2020. In: 2021 7th International Conference on Advanced Computing and Communication Systems (ICACCS), vol. 1, pp. 314–319 (2021). DOI 10.1109/ICACCS51430.2021.9441835
39. Takikawa, H., Nagayoshi, K.: Political polarization in social media: Analysis of the "twitter political field" in japan. In: 2017 IEEE International Conference on Big Data (Big Data), pp. 3143–3150. IEEE (2017)
40. Thelwall, M.: The heart and soul of the web? sentiment strength detection in the social web with sentistrength. In: Cyberemotions, pp. 119–134. Springer (2017)
41. Trabelsi, A., Zaïane, O.R.: Phaitv: A phrase author interaction topic viewpoint model for the summarization of reasons expressed by polarized stances. In: Proceedings of the International AAAI Conference on Web and Social Media, vol. 13, pp. 482–492 (2019)
42. Wong, F.M.F., Tan, C.W., Sen, S., Chiang, M.: Quantifying political leaning from tweets, retweets, and retweeters. IEEE transactions on knowledge and data engineering **28**(8), 2158–2172 (2016)
43. Wulf, V., Aal, K., Abu Kteish, I., Atam, M., Schubert, K., Rohde, M., Yerousis, G.P., Randall, D.: Fighting against the wall: Social media use by political activists in a palestinian village. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, pp. 1979–1988 (2013)