ORIGINAL ARTICLE

Social media analytics system for action inspection on social networks

Marco Mameli[1] · Marina Paolanti[1,3] · Christian Morbidoni[2] · Emanuele Frontoni[1,3] · Antonio Teti[2]

Abstract
Social networks are increasingly used for discussing all kinds of topics, including those related to politics, serving as a virtual arena. Consequently, analysing online conversations, for example, to predict election outcomes, is becoming a popular and challenging research area. On social networking sites, citizens express themselves spontaneously regarding political topics, often driven by specific events in social life. Real-time analysis of social media can provide valuable feedback and insights to both politicians and news agencies. In this paper, we discuss the design and implementation of a system for tracking and analysing social media. The SocMINT system provides an easy-to-use, visual dashboard to monitor the discussion on specific topics, to capture trends in communities and, by iteratively applying multidimensional data analysis and filtering, to deeply analyse posts and influencers. SocMINT aggregates data from multiple social sources and performs sentiment analysis on textual, visual and mixed content via a specifically designed neural network architecture. The system was applied in a real context by administrative staff of a political party to effectively analyse candidates' political communication on Facebook, Instagram and Twitter and the related online community reactions and discussion. In the paper, we report on this real-world case study, showing how the system meaningfully captures trends in public opinion, comparing the main KPIs provided by SocMINT with the outcomes of traditional polls.

Keywords Social media · Data visualisation · Sentiment analysis · Deep learning · Social decision support system

1 Introduction

Monitoring and measuring public opinion trends on social networks has emerged as a potential alternative to opinion polls due to its spontaneous nature and accessibility to a large number of people (Tsaytler et al. 2011). However, as an original and passive data source, social media data pose several analytical challenges, such as how to identify the target information from unstructured and unprompted data, how to quantify the highly qualitative visual and textual messages and how to ensure representative data of an examined topic (Gong et al. 2020).

In recent years, as the popularity of social media has increased, considerable research efforts have been made to predict elections outcomes based on sentiments expressed on social networks. The main objective has usually been to evaluate sentiment scores from the related social media posts as accurately as possible, for example, to predict the winning candidate (Liu et al. 2020). In fact, conducting a poll survey is typically very expensive and time-consuming. In contrast, social media offers active feedback since the users are available in real time. In this regard, politicians have expanded the number and types of actors potentially able to shape the political discourse, and they have set their agendas considering these kinds of environments. Social media platforms have become ubiquitous communication channels for candidates during their election campaigns. Facebook, Instagram, Twitter and other platforms enable candidates to reach out to voters directly, mobilise supporters and influence the public opinion. These important changes in political communication present strategic choices for the candidates.

Research on social media effects on elections has concentrated on data generated by communication conventions, such as retweets, mentions, likes or hashtags. Understanding the ways in which politicians adapt their message contents to the peculiarities of different platforms for political

✉ Marina Paolanti
marina.paolanti@unimc.it

1   VRAI Vision Robotics and Artificial Intelligence Lab, Department of Information Engineering (DII), Università Politecnica delle Marche, Via Brecce Bianche, Ancona, Italy

2   Università degli Studi "G. d'Annunzio" Chieti, Pescara, Italy

3   University of Macerata , Macerata, Italy

communication remains an ambitious goal (Cheung and Thadani 2012; Wang and Zhang 2012). After this customisation, the information is then disseminated to the public. This extra step requires additional time for validation, with a very short window for gaining consensus. Social media analytics systems offer a promising way to improve public awareness and action (Wu et al. 2007).

In this paper, a novel social media analytics system is presented. The proposed system, called SocMINT, provides the ability to collect data over different social networking platforms and offers users an interactive dashboard that supports multidimensional analysis and incorporates a deep learning-based method to predict social sentiment related to both images and textual data. A method for estimating the overall sentiment of a picture based on both visual and textual information is proposed, as previously done in Paolanti et al. (2017), and incorporated in the analytics system. The model uses a two-stage deep learning classifier. The first stage is for the image, while the second is for the text inside the image. In this way, we are able to include in the analysis content such as *MEMEs*, which are widely used on social platforms. The visual feature extractor is based on the VGG16 network architecture (Simonyan and Zisserman 2014) and is trained by fine-tuning the base model on ImageNet (Krizhevsky et al. 2012). The visual feature extractor is applied to the whole image, while the textual feature extractor detects and recognises text in images before performing sentiment analysis. The textual sentiment classification is based on BERT layers (Devlin et al. 2018) with the use of transfer learning of the classifier using a text sentiment dataset.

Finally, a real case study is reported, in which SocMINT is used by the staff of a political party to monitor and analyse the social activity and political communications during the campaign for the election of the president of the Marche region in Italy on 20–21 September 2020. Beside demonstrating the functionalities of the platform in analysing and gaining insights from the data, we show how the key performance indicators (KPIs) measured by SocMINT can be combined to reflect the opinion trends towards the two main candidates by benchmarking KPIs against the results of traditional polls.

The contribution of this paper is twofold:

- The design and implementation of a social data monitoring and analytics framework, incorporating cross-media sentiment analysis, graph analysis-based KPIs and multidimensional data exploration in a user-friendly web dashboard.
- The validation of the framework by its application to a real case study related to local political elections.

The paper is organised as follows: Sect. 2 includes a review and discussion of previous studies to provide the background and foundation of our study; Sect. 3 describes in detail the core models of our work and provides a detailed discussion of the sentiment analysis procedures; the final sections present the experimental results (Sect. 4). In conclusion, final remarks and future perspectives are reported in Sect. 5.

2 State of the art

A number of research studies prove that social networks can be successfully used to predict outcomes in many areas (Colladon et al. 2019; Daud et al. 2020). Predictive analysis with social media is an emerging tendency, which can be employed in numerous real-life applications (Lassen et al. 2017).

Prediction based on social media data is used by numerous companies, industries and organisations to enhance their business by predicting trends and behaviours (Kaplan and Haenlein 2010; Zafarani et al. 2014).

In Yu and Kak (2012), the authors analysed subjects including movie box, marketing, elections, information dissemination and microeconomics to make precise predictions using social media data. Yu and Kak classified predictors implemented in the literature (message and social network characteristics) along with their metrics (sentiment, time series, terminology, degree, density, centrality, structural hole). Finally, they specified the state-of-art prediction methods as regression methods, the Bayes classifier, the k-nearest neighbour classifier, artificial neural networks, decision trees and model-based prediction.

Asur and Huberman proposed a taxonomy of prediction models and provided the corresponding literature for four sub-fields: electoral results, stock market movement, product sales and influenza incidence. They also investigated difficulties and factors for forecasting modes in social media (Schoen et al. 2013)

In Lassen et al. (2017), the authors examined the differences between predictive and explanatory models along with the advantages of predictive models.

A model based on Thelwall's SentiStrength emotion detection algorithm was developed in Burnap et al. (2016). The authors used around 14 million Tweets to predict the results of the UK general elections.

In Karami et al. (2018), the authors presented a computational public opinion mining approach, collecting data from Twitter. In addition, they utilised machine learning classifiers and Linguistic Inquiry and Word Count (LIWC) and applied the Mallet implementation of latent Dirichlet allocation (LDA) to explore, detect and analyse economic issues in social media during the 2012 US presidential elections. Their methodology involved evaluating the difference between the number of positive topics and the number of negative topics, indicating that Obama had the advantage

over Romney, and it outperformed the Pew Research Center survey results.

Thirteen different variables were used in Isotalo et al. (2016), including Tweets, celebrity Tweets and celebrity sentiments, Twitter followers, Facebook page likes and Wikipedia traffic, for the 2016 US presidential elections. The research identified correlations between polls and Facebook page likes and between polls and Twitter. The authors applied linear regression analysis with data partitioning and scoring and concluded that machine learning models with linear regression could predict election outcomes with meaningful accuracy.

In Cameron et al. (2016), social media data were used for the prediction of election results in New Zealand. This work demonstrated that the numbers of Facebook friends and Twitter followers as explanatory variables were not good indicators, as only 16.7 and 5.4% of election winners were predicted correctly, respectively.

In MacDonald and Mao (2015), the authors used data from social media as well as from other sources, including Google trends, Wikipedia, polls and the news. Text mining methods and the vector auto-regressive (VAR) technique were combined in a framework to predict the results of the 2015 Scottish and UK elections, yielding impressive prediction accuracy. They correctly predicted the rank of the parties and often made very accurate predictions within decimals (i.e. the mean rate for the percentage of the Conservative party in Scotland was predicted to be 14.73% and the actual one was 14.90%).

In addition to predicting election outcomes, actors in the political scene need tools to perform in-depth descriptive analyses, for example, to understand public opinion about specific topics or to measure the influence of politicians, columnists or social media users in general. The ability to derive valuable insights from online communications and discussions, in turn, is important to support decision-making. Sentiment analysis is central to decision support systems involving social media analysis (Ducange et al. 2019; Gruhl et al. 2010).

The first example of using social analysis for political elections is Park et al. (2011), in which the authors analysed how social networks allow diplomats to establish direct communication with the population. A different approach is used in Park et al. (2011), where sentiment analysis is employed to understand the political views of people based on their comments on political news. In recent years, some political elections have been studied from a social perspective using sentiment analysis based on social network discussions; for example, Kušen and Strembeck (2018) used the Twitter date of the 2016 Austrian presidential elections and applied the sentiment and social analysis, with the aim of predicting elections winners. The same approach, once again based on Twitter, was used by Epp (2016) for the 2015 Canadian

election. The authors developed a randomised model based on sentiment analysis in an attempt to understand sentiments about the same argument outside of the social network environments, that is, in the real world. An additional example of social sentiment analysis applied on Twitter and focussed on Great Britain's referendum concerning a potential exit from the European community, also called Brexit, was reported by Agarwal et al. (2018). The authors analysed tweets by performing geospatial sentiment analysis based on event locations compared to geospatial tweet distributions for the particular events on a global level. They demonstrated that people mostly used keywords and hashtags during that time and showed how these words and symbols were being applied (i.e. which hashtag was used positively, which was used negatively and which indicated a neutral sentiment).

As social networking sites are becoming the primary means to communicate in nearly all contexts, effectively integrating different types of data analytic tools is an important issue in investigating and understanding public discussions on social and political topics. An extensive survey of social networking data analysis tools was conducted in Sapountzi and Psannis (2018), and challenges and future directions were highlighted. The authors concluded that there is a lack of tools that integrate data from different social networking sites and that, in general, data processing and analysis still requires several skills and different tools. Furthermore, they emphasised the need to integrate the topological network structure with content mining as well as the importance of relying on multi-modal sentiment analysis.

A framework to analyse online political communication was presented in Stieglitz and Dang-Xuan (2013) with the aim of providing guidelines for developing toolsets that encompass data collection, storage, monitoring and analysis of content from social media for political institutions.

An analytical framework was presented in Coletto et al. (2017) that explicitly targets polarised discussions in online social networks and applied to two case studies: online discussions about emigrants in Europe and about Brexit. This was relevant in this context, as most politics-related discussions, including those related to elections, are typically highly polarised, especially on social networking sites. The idea involved mixing multiple analysis dimensions (i.e. spatial, temporal and sentiment). However, social network analysis dimensions were missing from this framework. Brexit was also investigated by Grčar et al. (2017) in relation to Twitter massages. The stance towards or against Brexit was automatically detected by machine learning and compared with the outcomes of traditional polls, and influencers in both pro-Brexit and against-Brexit communities were detected by simply combining the amount of retweets and original tweets produced by users.

Most of the works focussing on political elections have used the Twitter platform as the primary data source; on the

Fig. 1 The SDSS architecture. The social media data are collected and stored in a database. The evaluation metrics are defined and saved inside the database as KPIs. A deep learning approach is then used for sentiment prediction. Finally, SocMINT, a web-based system, is designed to provide insights and performance perspectives
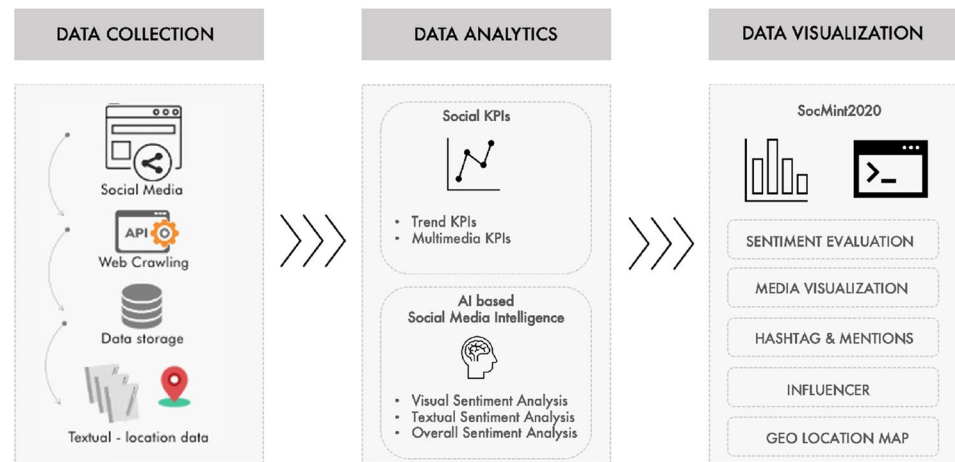
contrary, the analysis platform presented with this work used data from Facebook, Twitter and Instagram.

3 Data, measures management and sentiment analysis

The proposed system architecture has the main purpose of predicting the opinions of users on real data coming from social media in near real time.

The system workflow consists of three main phases, which are shown in Fig. 1 and detailed in the following subsections:

- Data Collection (Sect. 3.1),
- Data Storage (Sect. 3.2),
- Social Key Performance Indicators (Sect. 3.3),
- Sentiment Analysis (Sect. 3.4),
- SocMINT (Sect. 3.5).

3.1 Data collection

In our experiments, data are collected from three sources: Facebook, Instagram and Twitter.

The data collection module can be used with different types of sources: public/authenticated APIs, such as the open Twitter API, static data files (e.g. collected from paid services) or web crawlers/scrapers, to get content from unstructured HMTL data sources.

Users can collect live data about a topic from selected data sources by creating "data campaigns". A data campaign is defined by a query on hashtags and keywords, possibly consisting of AND and OR logic operators and a list of social media accounts. When a campaign is created, posts matching the query and/or from the specified accounts are incrementally collected into the corresponding dataset.

The data are preprocessed to extract the so-called Rawposts Information, which includes the text of the post, publication date, author, list of hashtags and user mentions found in the post as well as links to media in the post. For each post, we also store the number of likes or reactions, the number of reposts and geolocalisation if available. All distinct post authors are collected in a relational table along with relevant information, including standard metadata and KPIs (e.g. number of friends, followers, posts) and connections among authors who have interacted with the same content. The latter information will be used to derive advanced KPIs, as described later. Finally, media included in processed posts are collected along with information about source devices (mobile, computer, etc.) and ranked by the number of posts linking to them.

3.2 Data storage

A database management system (DBMS) based on representation through documents is used for the storage of the collected data. This class of DBMS, which is part of NoSQL, includes several types. Among these, MongoDB is selected here based on the works of Kanoje et al. (2015) and Abramova and Bernardino (2013), as it is considered the best DBMS for managing social network data. The database structure has three main fields: user, rawposts and posts. The first stores the information concerning the user, which can be extracted from the social network. Privacy is respected because only publicly shared information is accessed. Rawposts is a collection that contains the information extracted from the social network without any manipulation. It includes all the data provided by Facebook or Instagram through a post or by Twitter through a tweet. Finally, posts are a collection of data that will be used for the visualisation in SocMINT.

3.3 Social key performance indicators

In this section, we discuss the social KPIs adopted by the proposed system to measure the performances of social media posts and users. We distinguish between trend and multimedia KPIs.

3.3.1 Trend KPIs

These types of indicators are used to track trends that characterise a phenomenon over time. Social media trends can vary since there are different "phenomena" to be analysed. The trend KPIs are engagement, number of published posts, degree centrality, closeness centrality and entropy.

Number of Published Posts The number of published posts containing a set of keywords or hashtags provides a simple way to evaluate the interest generated by a topic (e.g. a set of keywords of hashtags) as well as the interest of a given user in the topic, when posts from a single author are considered.

Engagement Engagement is a common measure that reflects the interest expressed by users towards a given content by means of direct interaction with it Muñoz-Expósito et al. (2017). It can be simply defined as follows:

$$\text{Engagement} = n \, \text{likes} + n \, \text{comments} + n \, \text{shares}. \tag{1}$$

Engagement can be measured for a single post or averaged over the number of published posts by the same author to measure her ability to engage people.

All three tracked social media platforms are covered in Eq. 1, where n likes is the number of likes obtained by a post on Facebook plus the number of *hearts* obtained on Twitter and Instagram, n comments is the number of comments on the post and n shares is the number of shares and retweets.

Centrality Centrality measures based on graph theory (Peng et al. 2018; Rabade et al. 2014) are used to evaluate the role of the user during the exchange of posts and/or comments in a social network. In particular, two metrics are selected for the proposed system: degree centrality and closeness centrality.

Degree centrality takes into account the number of neighbours of a node in the network (Rabade et al. 2014). It is used to measure the potential ability of a node to disseminate information in the network. In case of social media, the degree centrality of a user can be measured by counting the user's number of friends (Facebook), followers and following (Twitter and Instagram), while the degree centrality of a community group or page (both open and private) can by evaluated by counting subscribers and users who have subscribed, showing interest in the contents of the page. Degree centrality is defined as follows:

$$C_D(i) = \sum_{j=1}^{n} a_{ij} \tag{2}$$

where a_{ij} is 1 if the connection from node i to node j exists, and 0 otherwise.

Closeness centrality, instead, takes into account the distance between nodes in the network, measuring, for each node, the average distance from the connected nodes. Provided a distance metric d, the closeness centrality can be defined as follows:

$$C_C(i) = \frac{1}{\sum_{j=1}^{n} d(i,j)} \tag{3}$$

where $d(i, j)$ represents the distance between nodes i and j.

Closeness centrality can be used to measure the influence of users in a community, as it provides an efficient measure of the dissemination of information within a network (Rabade et al. 2014), following the assumption that the higher the proximity centrality value of a node, the better the positioning of this node is in the social network.

Different distance metrics can be used to calculate the KPI. In our system, we assume there is a connection between user i and user j if user j engaged with a post (or share) made by user i. The weight $W_{i,j}$ of the connection is the interaction rate, provided by the number of posts from user i which j interacted with divided by the total number of posts from user i. We use the inverse of $W_{i,j}$ to measure distance between the two nodes.

The centrality considers only neighbouring nodes that represent adequate but not complete information. This limit has been imposed by the platforms because if the profiles are not public it is not possible to access information such as the number of posts and friends.

Entropy Entropy is a concept borrowed from information theory, which, in general, measures how much information is carried by an event, assuming that the more uncertain and random an event is, the more information it potentially contains. He et al. (2013) introduced transfer entropy to measure the amount of information exchanged on a social network. In Peng et al. (2015), it was then used and adapted to model a mobile network in which users exchange SMSs, which led the authors to define *friend entropy* as the measurement of users' influence within the mobile social network. Further metrics presented by Peng et al. (2015) apply entropy to the interaction frequency among users, measuring the number of interactions between users in the mobile network.

Friend entropy can be easily applied to social networks, as it is based on the number of node connections in the graph, by considering the number of friends (Facebook) and the number of followers (Twitter and Instagram)

Fig. 2 Sentiment analysis work-flow. The rawpost is used as input. In the top of the schema, the image is used inside the image classification network. At the same time, the image is the input of the text extraction. The text of the post, along with its extraction, becomes the input of the text classification
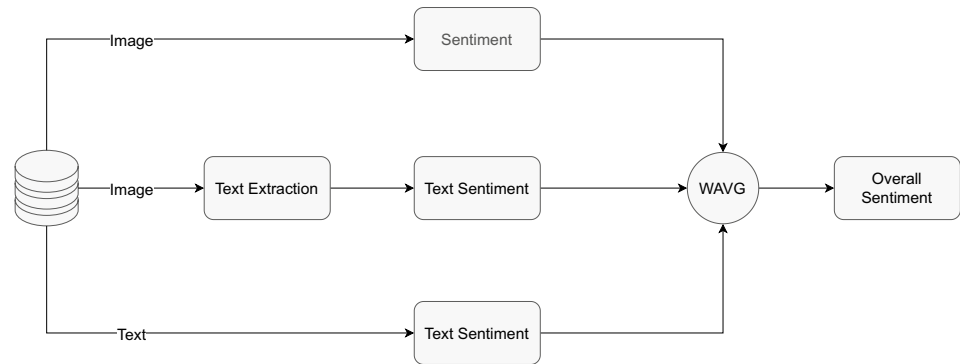
Defining $N_i(t)$ as the number of friends and followers of user i at time t, we can define friend entropy as in Eq. 4, as done in Peng et al. (2015).

$$I_i^f(t) = -\frac{1}{N_i(t) + \theta} \log_2 \frac{1}{N_i(t) + \theta} \qquad (4)$$

where θ is a parameter equal to 1 to take into account the possibility that there are users who do not have friends.

In Peng et al. (2015), frequency interaction entropy was used on a mobile network, where direct interactions via SMS were tracked. On social media, however, while direct messages are sometimes allowed (e.g. on twitter via the mention @), the public nature of the communication results in most interactions being content based. As in the case of centrality, we therefore consider engagement actions with the content.

Therefore, defining $C_j(t)$ as the number of interactions obtained by post j at time t and N_p as the number of posts of user i, we can define the interaction frequency entropy as in Eq. (5).

$$I_i^c(t) = -\frac{1}{\sum_{j=1}^{N_p} C_j(t) + \lambda} \log_2 \frac{1}{\sum_{j=1}^{N_p} C_j(t) + \lambda}. \qquad (5)$$

Finally, the entropy measure used for the trend KPI is given by a combination of the described measures, as reported in Eq. (6).

$$I_i(t) = \alpha I_i^f(t) + \beta I_i^c(t) \qquad (6)$$

where $\alpha + \beta = 1$, and they denote the importance assigned to each entropy measure. While such weights could certainly be tuned in a given context, in our experiments we simply assign the same weight to the two entropy measures.

3.3.2 Geospatial KPI

Geospatial KPIs allow real-time operational views of people, services, assets and events. In particular, these make the monitoring system practical and effective. Geography information is the key to understanding real-world big data

and integrating it into users' daily activities and operations to provide important insights. This is because geographically based data specify where things occur and provide contextual knowledge for social media data.

3.3.3 Multimedia KPI

The evaluation of the overall sentiment for a post will be the combination, or more precisely, the weighted average of the individual sentiment values obtained for the text and image linked to the post. The associated weights will be selected by the user to allow customisation. For both of these sentiments, the same values are analysed: 0 neutral, 1 negative, 2 positive. The sentiment analysis of an image consists of assigning a label, among those reported above, to an input image. This task requires the description of the image by means of characteristics that identify its content. Meanwhile, textual analysis is a complex task for examining linguistic differences, the length of the text, the context, etc. In addition, it is possible that the text published on social networks contains sarcasm. Based on these considerations, for SDSS it is crucial to use a textual analysis model that considers the entire text of the post as an input.

3.4 Sentiment analysis

The approach presented by Paolanti et al. (2017) (i.e. the combination of the visual and textual features) has been used and extended for development of the proposed SDSS. The framework for joint visual and textual analysis, as well as the novel data set used for evaluation, comprises three main components: the two sentiment classifiers (i.e. visual and textual sentiment) and the overall sentiment evaluation, as shown in Fig. 2. This framework is the basis for the evaluation of the multimedia KPIs previously described.

Sentiment classification considers three different integer values: 0 for a neutral sentiment, 1 for a negative sentiment and 2 for a positive sentiment, assigned by the user. The visual feature extractor is based on the VGG16 network architecture (Simonyan and Zisserman 2014), and it is

trained by fine-tuning a model pretrained on the ImageNet dataset (Krizhevsky et al. 2012). VGG16 has the best performance for image classification and offers the best trade-off between inference time and memory usage (Bianco et al. 2018). Transfer learning (Hussain et al. 2018) is applied by replacing the top layer of the network, and the fully connected layer classifies the input image in order to bring the number of possible output classes from 1000 to three. The transfer learning allows maintaining high precision in the extraction layers and training only the fully connected classification layer, with the aim of understanding what can or cannot be positive.

The textual feature extraction from images is performed using a two-step workflow based on the Character Region Awareness for Text Detection (CRAFT) designed by Baek et al. (2019), and the output is used as input for the spatial attention residual network (STAR-Net) of Liu et al. (2016). In particular, the CRAFT network takes as input the image of the post and extracts a list of the bounding box coordinates where there is text in the image. If the list of the bounding box is empty, the image is considered an image without text. In an image that has text (i.e. the list of bounding box is not empty), the bounding box coordinates are used to crop and extract the text section of the image. This cropped image with text inside becomes the input of the next STAR-Net, the network that extracts the text in the image. The final stage of the sentiment classification includes the overall sentiment analysis, with a weighted average of the single sentiment that is obtained from the image and text analysis. The weight associated to the single sentiment value is chosen by a communication expert, who evaluates the data and content and decides to give more importance to the text content of a post, previously called a captcha. For this reason, the multimedia KPI assigns greater importance (with a numerical value of 2) to the text of a post, a single importance of 1 to the image sentiment and 0.75 to the text sentiment extracted from an image.

Textual sentiment classification is crucial to obtain an overall sentiment of the post. The association of sentiment with text information is challenging task due to different languages. In addition, in a social context, the association of the sentiment with a text extracted from a post or an image can become more complex if it contains sarcasm. Accordingly, for the proposed SDSS system it is important to use a model that considers the entire post and also understands the context of the text input. To make this possible, natural language processing (NLP) is used with the application of the machine learning techniques. As noted by Zhang et al. (2018), NLP research has produced numerous results related to applying sentiment analysis based on machine learning methodologies, which use hand-crafted features with the application of supervised or unsupervised algorithms, as mentioned, for example, in Pang and Lee (2008), Liu (2012), Liu (2015).

In recent years, the use of deep neural networks has led to greater application of convolution neural networks, as in Johnson and Zhang (2014); Santos and Gatti (2014), or long short-term memory (LSTM), as in Katiyar and Cardie (2016); Qian et al. (2016). Even better results have been achieved using the attention model (Yang et al. 2017). Considering the importance of contextual relations with words for the system realised in this framework, and due to the demonstrated effectiveness of the transformer value using a trained neural network as the basis for a new purpose-specific model in computer vision research, an attention model is used here, specifically the Bidirectional Encored Representation from Transformer (BERT) (Devlin et al. 2018). The goal of BERT is to generate a language representation model based on a tokenised text input sequence as suggested by Rai and Borah (2020), for the NLP task. A pre-trained BERT layer, publicly available on tensorflow-hub,[1] is incorporated in our neural network model. The weights of the layer were learned from a multi-language dataset (not publicly available). We then fine-tuned the model by training it on a portion of our dataset in order to effectively apply sentiment prediction in our application domain.

Sentiment analysis applied in a combined way to text and images allows us to incorporate messages conveyed by visual memes, in which a text message in often incorporated within images. The result of the sentiment is stored individually for the text, the text extracted from the image and the image, as illustrated in Fig. 2. The predictions are then combined by calculating the weighted average of the three contributions. The value range is divided into three equal parts in order to assign a net sentiment value: 0.68 is used as the threshold for positive sentiment; a value between 0.34 and 0.67 indicates a neutral sentiment; and a sentiment is considered negative below 0.34.

3.4.1 Dataset for forecasting the elections in the Marche Region

The dataset built for the proposed SDSS includes content from Facebook, Instagram and Twitter. There are 181259 posts in total, of which 7.65% have an associated image (10.57% from Facebook and Instagram and the remaining part from Twitter). In order to obtain the ground truth necessary for training and evaluating the classifier, the post sentiments were manually annotated by three human annotators, who are students of Università Politecnica delle Marche. In cases in which the labelling resulted in, for one or more elements of the dataset, three different labels (which then made it impossible to select a final label), a fourth student intervened who was assigned only the

---

[1] https://www.tensorflow.org/hub/tutorials/bert_experts.
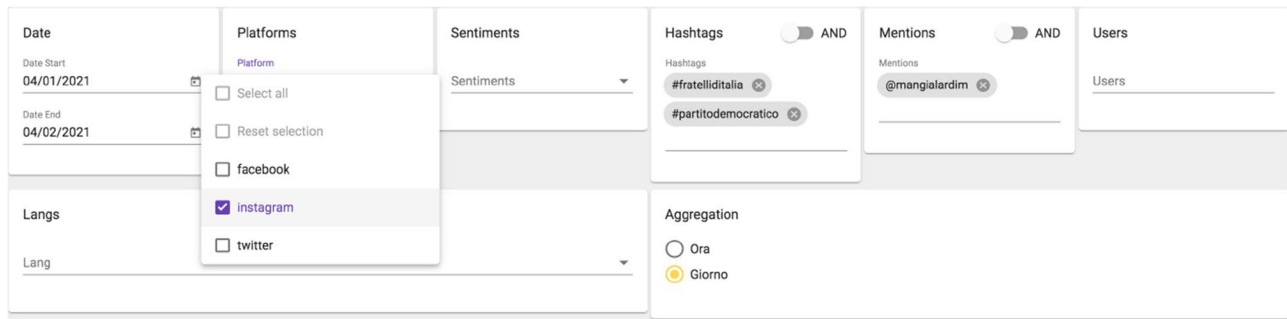
Fig. 3 SocMINT. Filters

particular cases to be analysed and, at the end of this analysis, selected the label with a higher cardinality weighting. This labelling procedure required a week of work by the first three students. Then, in an automatic way, the entries with three discordant labels were identified. These 100 entries required a further three hours for labelling. Three distinct annotations are made for each post, targeting their visual, textual and overall sentiments. After a month of text data, the BERT layer was fine-tuned on the new data to make the prediction of the sentiment of the text more precise. The dataset was split using the standard subdivision percentage: 20% for the test and the remaining 80% of the dataset for training, which was subdivided into training and validation sets, with the same percentage of the test dedicated to validation. This means that 5000 posts were used for the test and 20,000 for the validation and training set globally. This set-up choice was made based on the network architecture and the dataset itself. The proportions were selected based on the size and type of our data. In fact, it is important that the data that are being used to learn and those being used to validate or test the model have statistical distributions that are as similar as possible.

3.5 The SocMINT dashboard

The web-based system described in this paper is called SocMINT, and it is designed to visualise social media statistics and evaluations in real time. The main SocMINT visualisation and analysis dashboard, based on KPIs, metrics and raw data, is populated from a dynamic input, which can be updated by the user at any time, and is composed of the following:

- A list of keywords (e.g. Facebook with particular reference to pages);
- A list of hashtags (e.g. Twitter/Instagram);
- A list of social platforms accounts/handles (e.g. Twitter/Instagram).

The dashboard supports time-dimension analysis by providing dynamic time window selection, which, along with other filtering functionalities, can be used to segment the underlying data and focus on specific events. Different levels of aggregation granularity are supported, namely hours, days and weeks. In Fig. 3, the main filters provided by the dashboard are shown. By setting filter values, all the dashboard graphics and views are updated accordingly, thus enabling a quick look at specific data sources, events, hashtags and mentions.

SocMINT includes five main views: sentiment, multimedia, hashtag, influencers and map of location data.

The lending view is the based on post sentiments and is illustrated in Fig. 4. On the right side of the view, there is a list of posts that can be sorted by date, username, engagement and platform, and it can also be downloaded in CSV format. On the left, three graphics are shown, starting from the top:

- The volume of the posts over time, selected by the current filter values (red curve), and the overall engagement of posts for each time point (the yellow bars).
- The overall post sentiment results as a bar graph, where red bars represent a negative sentiment, green bars a positive sentiment and grey bars a negative sentiment.
- Finally, pie charts provide a quick comparison between the amount of posts per platform and the overall engagement by platform.

Figure 5 shows the page dedicated to multimedia analysis. On this page, there are graphs with statistics about the media and the images accompanied with the post, which can also be sorted by date and engagement.

Figure 6 depicts the hashtag analysis page. A quick-look summary of the hashtags and users' mentions used in the selected posts is reported via bubble charts, where more frequent terms correspond to bigger bubbles. The graphs are located on the right alongside the post number and sentiment graphs.

Fig. 4 SocMINT. Sentiment analysis visualisation

The influencer analysis page is displayed in Fig. 7. The two histograms show centrality and entropy for each post author. The top one (titled INFLUENCERS AVG) shows the closeness centrality normalised using degree centrality. This normalisation is used to give information about the number of interactions normalised on the number of users that follow the creator of a post. The second chart (INFLUENCERS TOTAL ENGAGEMENT) is the result of entropy, which allows obtaining an overall measure of the engagement obtained by posts of a user.

The last page is for the geolocalisation of posts. The map shows the number of posts tagged with the corresponding position (only for georeferenced posts) (Fig. 8).

This view can provide additional insights about the attention for the monitored event over the territory.

4 Results and discussion: Marche Region election case study

The data collection system and the dashboard presented in Sect. 3 were used to monitor and study a political case within Italy, the regional administrative elections in Marche, by focussing on the analysis of the social campaigns of the two main contenders, Francesco Acquaroli for "Fratelli d'Italia" (FDI) and Maurizio Mangialardi for the "Partito Democratico" (PD). As stated previously, this is not the first time that social analysis has been used to follow election campaigns, but unlike previous analyses, SDSS uses data from three different sources. Furthermore, unlike the previous works, a comparison of the data
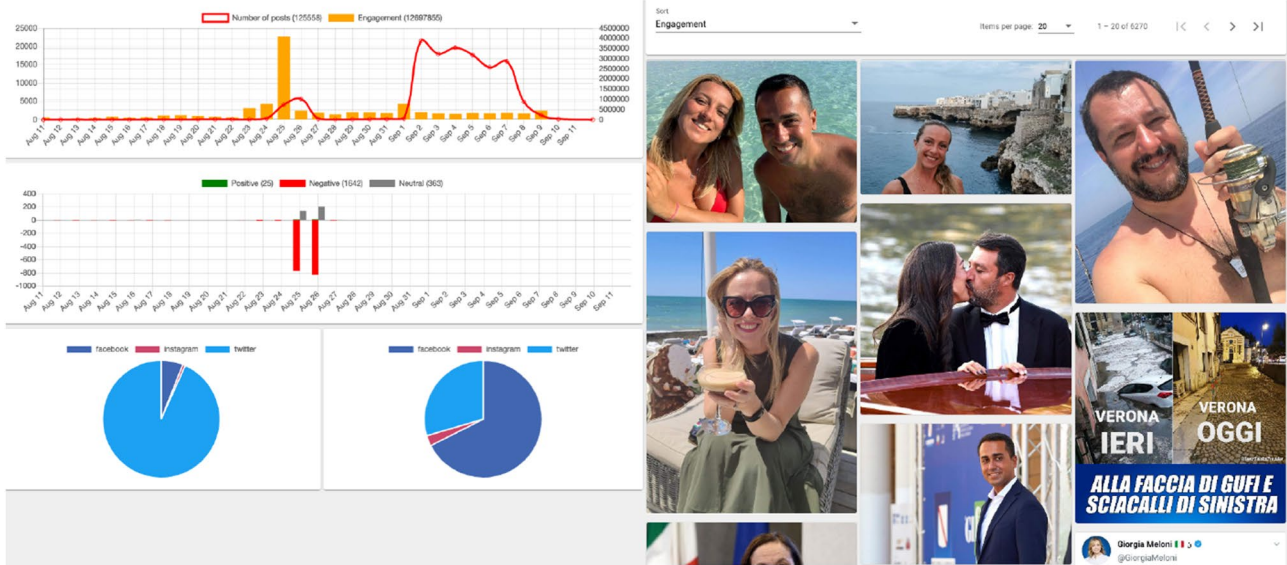
Fig. 5 Multimedia data analysis and visualisation
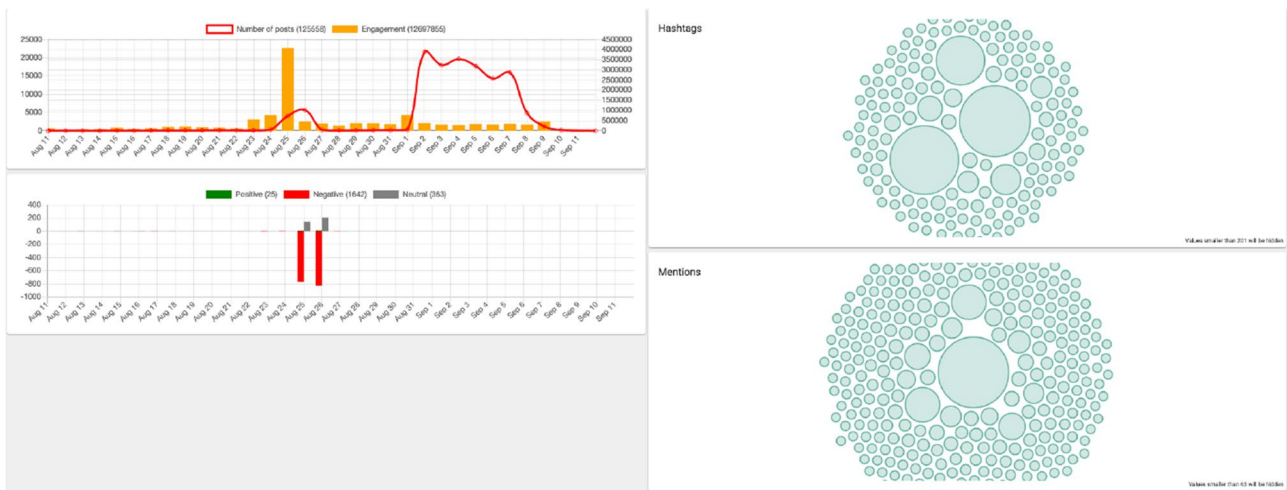

Fig. 6 Hashtag analysis and statistics

obtained from the dashboard is carried out with data from traditional surveys to determine whether SDSS can be used in real cases to drive social campaigns. The analysis is based on data collected by SDSS from 1 July 2020 until 27 September 2020. The beginning of the collection also coincides with the possibility of being able to hold and participate in public events due to the COVID-19 restrictions. The data collection date was chosen because it was exactly one week away from election day, and it was useful for studying the consensus of the winning candidates. The analysed dataset was collected with SDSS using the following hashtags: #acquaroli, #mangialardi, #ricostruiamolemarche, #marche, #regionemarche, #lega,

#fratelliditalia, #forzaitalia, #salvini, #meloni, #tajani, #renzi, #conte, #zingaretti, #berlusconi, #lemarcheintesta, #insiememarche. Such hashtags are related to political leaders in Italy from the main parties. These hashtags are also included in the query without the "#" symbol as keywords. We compose the query using AND and OR logical operators. More precisely, the hashtags #marche, #regionemarche, #ricostruiamolemarche, #acquaroli and #mangialardi are concatenated by the OR operator and by an AND statement with the remaining search terms. This allows us to exclude posts that are not correlated to the study case.

Using these hashtags to collect data resulted in a total of 181k posts, of which 7.65% have linked images and 10.57%
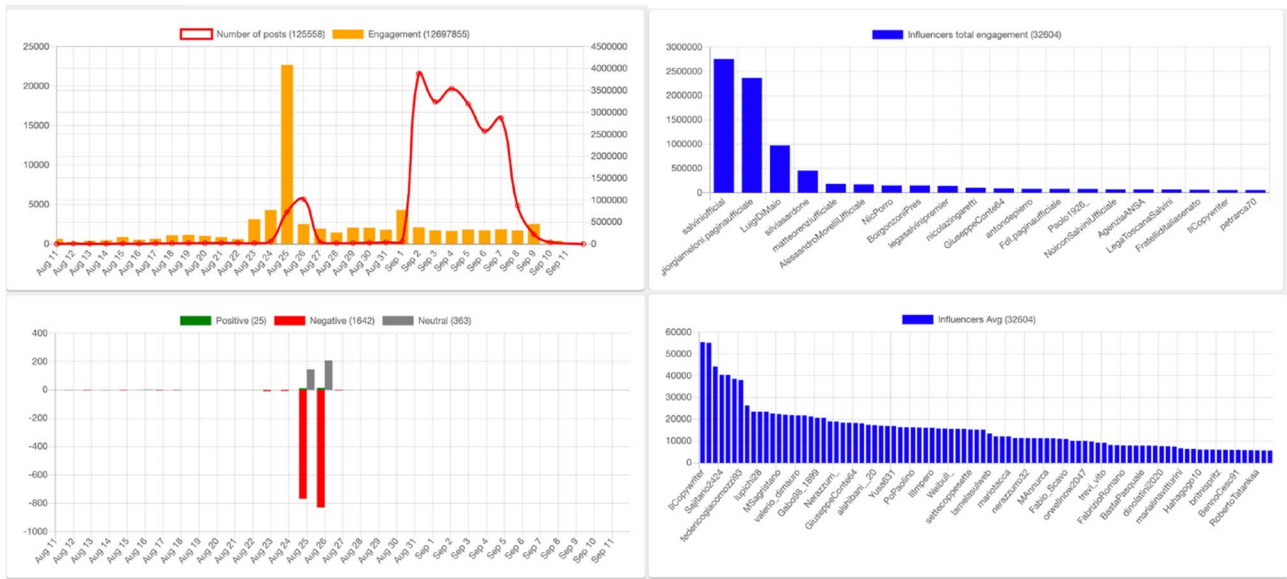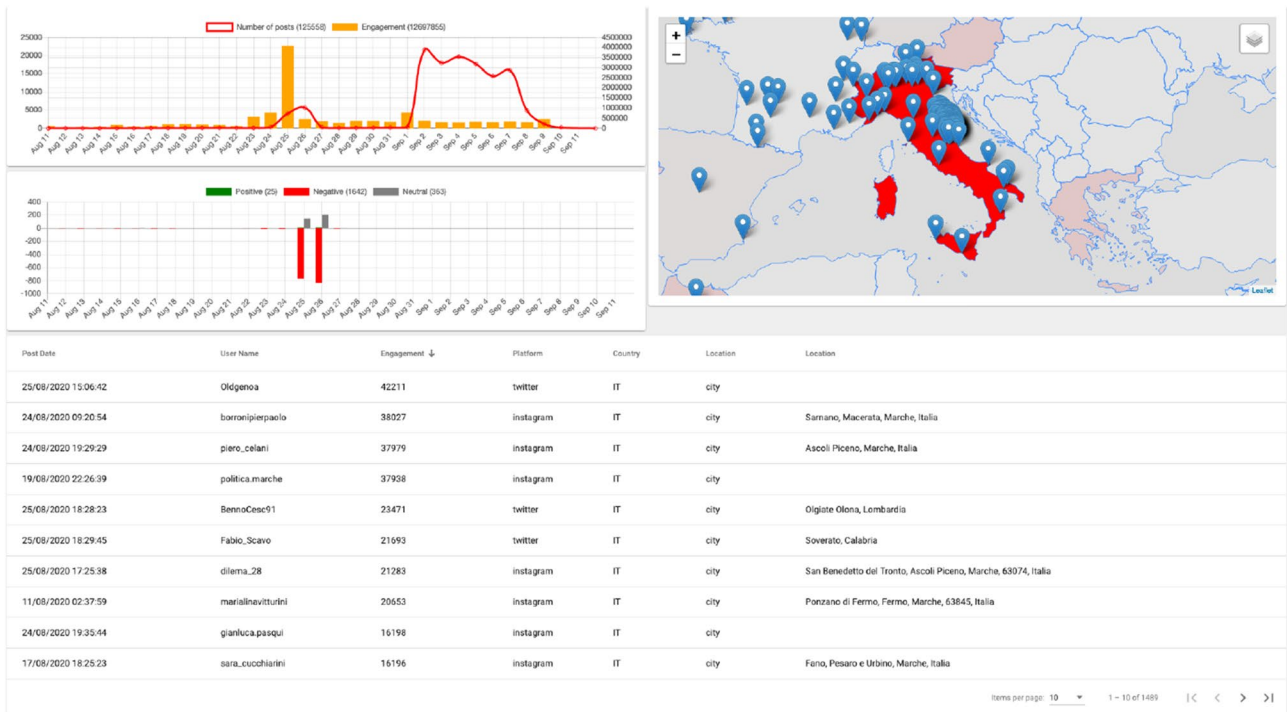
Fig. 7 Influencer analysis on SocMINT

Fig. 8 Map for the geolocalisation of the post

are from Facebook. Table 1 shows a summary of the post numbers by platform.

We used the posts collected during the first month to tune the sentiment classification module. In total, 25000 posts were labelled with sentiment annotation and used to further train the neural network. The labelling was performed as described in Sect. 3.4.1. To ensure objective labelling, the

purpose of the request was not explained to the students. They were asked to express a feeling of negative, neutral or positive for both the text and the image and, if present, for the text within the image as well.

The data collection started on July 1. For the whole month of July, as can also be seen in the graph extracted from SocMINT, the posts are 2.24% of the total, the greatest

concentration, and therefore, the greatest campaign activity starts in August.

In the remainder of this section, the analyses and comparisons will be reported, starting from 28 August, to compare the progress of the campaign with the trends of the classic political polls.

The analysis reported here compares the results of classical methods based on surveys and those provided by SocMINT. In Table 2, the results of both analyses are compared for a single week. The table reports three different metrics extracted from SocMINT: Personal Social Interest, Sentiment and Shared Social Interest. In particular, the Personal Social Interest column reports the overall percentage of interest by the community towards the candidate in the period of time under analysis. This value is obtained using the number of interactions per single post of the candidate. This analysis is accompanied by the average estimated sentiment of all the posts mentioning the two candidates, providing insight on the feelings of the community towards the candidates. The Shared Social column is probably the best candidate to reflect the traditional polls. The values in this column depend both on how the community interacts directly by sharing its content and on the interactions obtained by the sharers, so they make use of entropy and centrality to understand how interesting the topics are, considering the activities of people close to the candidate or those who are marginally part of his network.

In addition, in our case, as expected, the majority of the geotagged posts come from the Marche region, although a considerable amount come from north Italy and a few from other countries. This suggests that the population of north Italy is more active in the discussion, possibly because the support for right-wing candidates is stronger in the northern regions of the country. We notice, however, that the depicted picture is incomplete since not all of the posts are geotagged.

The comparison is shown in Table 2. Even with an approximation to excess, the analysis beginning with the dashboard allows monitoring the progress of an election campaign starting from a social campaign.

Besides partially reflecting poll results, thus providing help in evaluating the electors' opinions at a given time, the system allows performing a more in-depth analysis with respect to different aspects of a campaign.

In the case of a political campaign, salient events can be individuated by looking at peaks in users' engagement over time.

Some interesting insights can be obtained by analysing some moments of the monitored weeks. For example, in week 1, the Acquaroli committee, due to a typo on a press release, received considerable media attention, which it used to its advantage with a public apology, as reported in Table 3. It should be noted that the potential transmission on social media by the FDI leader Giorgia Meloni also helped.

This shows, also based on the position of the FDI parent company, the influence the network of people associated with an interest can have in achieving a positive result.

Moreover, the support provided by the presidents of the national parties to which the candidates belong may be influential. As can be seen in the bottom graph in Fig. 9, Giorgia Meloni, the president of FDI, is in the second position while Nicola Zingaretti, secretary of the PD, is at the bottom of the list of influencers. This suggests that the higher support from FDI on social media to the local candidate might have played an important role in the overall online community opinion (Fig. 10).

Other specific events, such as debates and confrontations between candidates, were analysed as well, aiming to understand how they were perceived by the online social community. The results are reported in tabular form in Table 3, and some examples of posts are extracted thanks to SocMINT.

For the other events, the SocMINT filters used are as follows:

- First public debate: a date filter for which we have used the date of "4/09/2020",
- Conflagration: a filter based on the date "15/09/2020" and the keyword "esplosione, porto",
- Voting day and vote result: are obtained by using the date filter, inserting the vote date and the date of the day after the result of publication (Table 4).

The second week consisted of a series of public meetings of the candidates with various business and tourism associations. In addition, a direct debate between the two political forces was held in order to discuss how they intended to advance their region. From these comparisons, both directly with the voters and with their direct opponent, the two candidates obtained a strong personal interest rank, indicating their ability to positively influence their voters. However, once again, despite the overestimate by SocMINT, the increased interest and confidence in the FDI candidate was confirmed, as shown in Table 5.

In the third week, the candidates commented on news events of regional importance and participated in a televised debate that received considerable interest from voters. In particular, the debate on labour issues allowed the PD candidate to capture the attention of the voters Tables 6 and 7.

In the week after the voting ended, the FDI candidate celebrated his victory. The indicators provided by SocMINT during this week suggest that the engagement of FDI voters was high, while low interest was shown by the PD voters in terms of discussing or commenting on the electoral defeat.
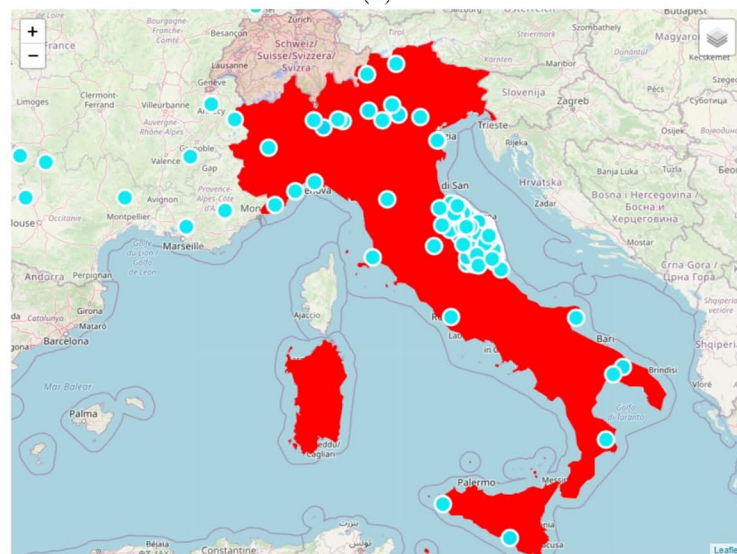
Fig. 9 Number of post differences between July and August and maps visualisation
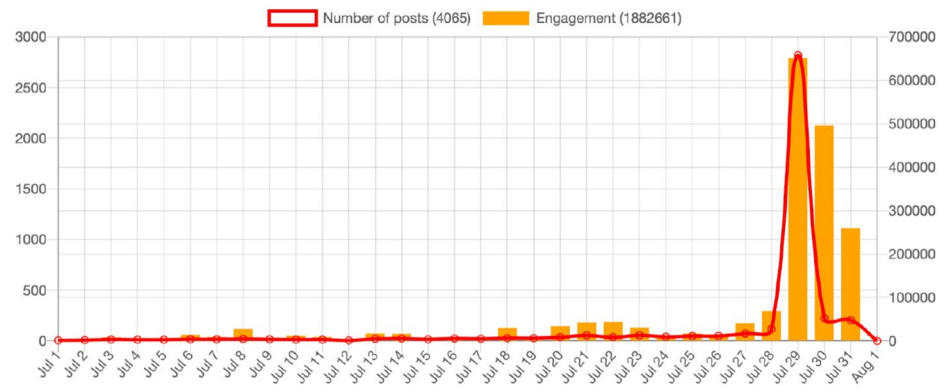
(a)
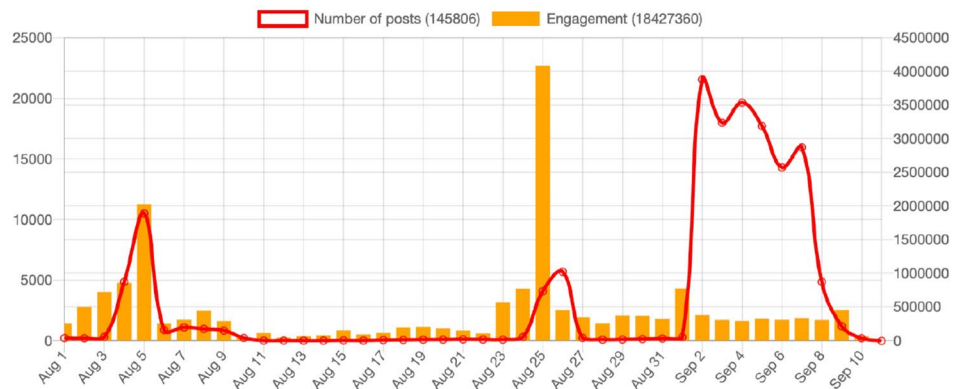
(b)

(c)

5 Conclusion

Extracting and visualising opinion polls surveying users of online social networks allows forecasting the results of an election. In this paper, a novel SDSS is presented, aiming to evaluate public opinion about a topic or phenomenon. The analysis is focused on political communication on Facebook, Instagram and Twitter by candidates during the

Table 1 The number of posts for the social platform

Dataset	Number of posts
Total	181,259
Facebook	17,691
Instagram	1466
Twitter	162,433

election campaign in the Marche region of Italy on 20–21 September 2020. The proposed SDSS utilises deep learning algorithms to predict social sentiment from both pictures and textual contents. Moreover, a web-based system called SocMINT has been built to visualise statistics and evaluations in real time.

The results confirm that our SDSS is suitable and effective for monitoring and evaluating users' opinions. Moreover, it is possible to examine the effects of multiple social media platforms along with the effects of conventional media, investigating their relative importance and their inter-relatedness.

Future works will be devoted to increasing the amount of statistics available. New metrics and new variables will be considered. Further investigations are needed to improve and better integrate the deep learning algorithms with more complex architectures to improve performance. Moreover, for sentiment analysis, the framework proposed by He et al. (2020) could be integrated for the analysis of images. It would ensure that new images samples extracted from social networks do not require new labelling procedures and offline training even in different scenarios. For text analysis, the *Learn#* approach presented by Shan et al. (2020) will be examined. These approaches will allow adapting the system to new events, improving the generalisability of the

Table 2 Week review summary

Week	Candidate	Personal social interest (%)	Sentiment	Shared social interest (%)	Polls (%)
First week	Acquaroli FDI	65.2	Positive	97.45	47
	Mangialardi PD	34.8	Negative	2.55	39
Second week	Acquaroli FDI	73.04	Positive	77.05	51.5
	Mangialardi PD	26.96	Positive	22.5	36
Third week	Acquaroli FDI	52.40	Positive	62.05	51.5
	Mangialardi PD	47.59	Positive	37.95	36

Fig. 10 **a** A histogram representation of the entropy measure for users, where Giorgia Meloni is in the second position. **b** The centrality of the user and the leader of FI is in the second position
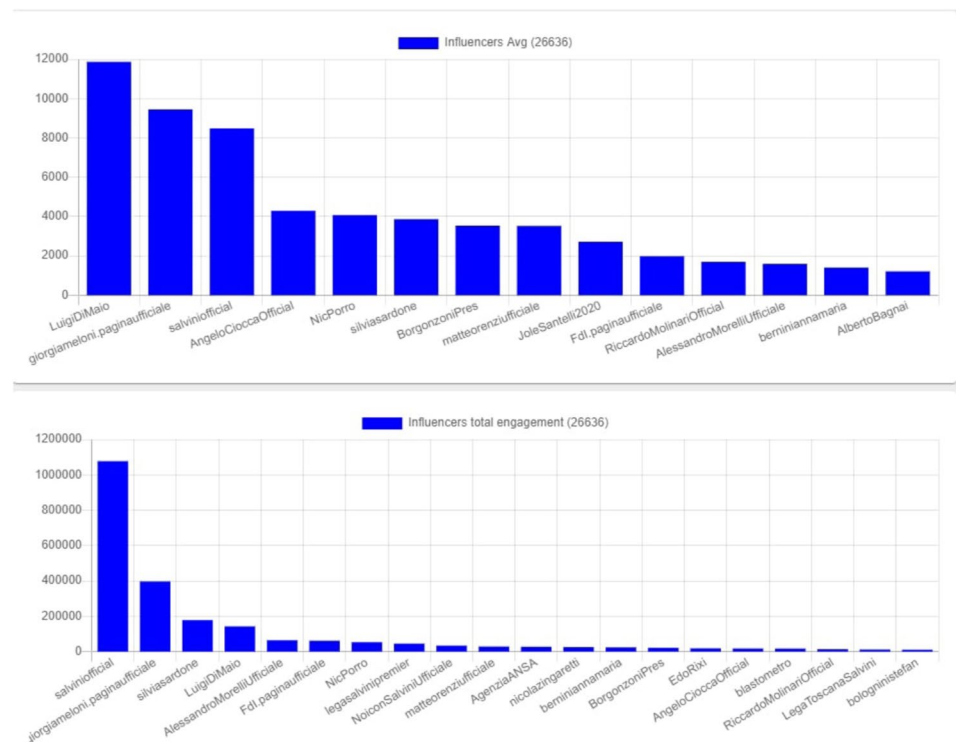
Table 3 Event analysis summary

Event	Candidate	Shared social interest (%)	Sentiment
Misprint	Acquaroli FDI	51	Positive
	Mangialardi PD	49	Negative
First public debat	Acquaroli FDI	68.3	Positive
	Mangialardi PD	31.6	Positive
Conflagration	Acquaroli FDI	31.5	Positive
	Mangialardi PD	68.5	Positive
Voting day	Acquaroli FDI	56.53	Positive
	Mangialardi PD	43.47	Positive
Vote result	Acquaroli FDI	90.81	Positive
	Mangialardi PD	9.19	Positive

Table 4 The percentage of the interest generated by the candidate compared with the official opinion polls

Candidate	Personal social interest (%)	Sentiment	Shared social interest (%)	Polls (%)
Acquaroli FDI	65.2	Positive	97.45	47
Mangialardi PD	34.8	Negative	2.55	39

Table 5 The interest generated by the candidate compared with the official opinion polls

Candidate	Personal social interest (%)	Sentiment	Shared social interest (%)	Polls (%)
Acquaroli FDI	73.04	Positive	77.05	51.5
Mangialardi	26.96	Positive	22.5	36

Table 6 The interest generated by the candidates compared with the official opinion polls

Candidate	Personal social interest (%)	Sentiment	Shared social interest (%)	Polls (%)
Acquaroli FDI	52.40	Positive	62.05	51.5
Mangialardi PD	47.59	Positive	37.95	36

Table 7 The two final moments of the elections: the first (Voting day) is the shared social interest of the social community in the fact that a candidate had expressed their vote, and the second (Result) is the percentage of people, based on messages of thanks and support, who were positive and happy about the results

Candidate	Voting day (%)	Result (%)
Acquaroli FDI	56.53	90.81
Mangialardi PD	43.47	9.19

Both of the percentages indicate high positive sentiment

system. We will also integrate automatic identification of non-genuine accounts.

References

Abramova V, Bernardino J (2013) NoSQL databases: MongoDB vs cassandra. In: Proceedings of the international C* conference on computer science and software engineering, pp 14–22

Agarwal A, Singh R, Toshniwal D (2018) Geospatial sentiment analysis using twitter data for UK-EU referendum. J Inf Optim Sci 39(1):303–317. https://doi.org/10.1080/02522667.2017.1374735

Baek Y, Lee B, Han D, Yun S, Lee H (2019) Character region awareness for text detection. In: Proceedings of the IEEE conference on computer vision and pattern recognition, pp 9365–9374

Bianco S, Cadene R, Celona L, Napoletano P (2018) Benchmark analysis of representative deep neural network architectures. IEEE Access 6:64270–64277

Burnap P, Gibson R, Sloan L, Southern R, Williams M (2016) 140 characters to victory?: Using twitter to predict the uk 2015 general election. Elector Stud 41:230–233

Cameron MP, Barrett P, Stewardson B (2016) Can social media predict election results? evidence from New Zealand. J Polit Mark 15(4):416–432

Cheung CM, Thadani DR (2012) The impact of electronic word-of-mouth communication: a literature analysis and integrative model. Decis Support Syst 54(1):461–470

Coletto M, Esuli A, Lucchese C, Muntean C, Nardini F, Perego R, Renso C (2017) Perception of social phenomena through the multidimensional analysis of online social networks. Online Soc Netw Media 1:14–32. https://doi.org/10.1016/j.osnem.2017.03.001

Colladon AF, Guardabascio B, Innarella R (2019) Using social network and semantic analysis to analyze online travel forums and forecast tourism demand. Decis Support Syst 123:113075

Daud NN, Ab Hamid SH, Saadoon M, Sahran F, Anuar NB (2020) Applications of link prediction in social networks: a review. J Netw Comput Appl 166:102716

Devlin J, Chang MW, Lee K, Toutanova K (2018) Bert: pre-training of deep bidirectional transformers for language understanding. Preprint arXiv:1810.04805

Ducange P, Fazzolari M, Petrocchi M, Vecchio M (2019) An effective decision support system for social media listening based on cross-source sentiment analysis models. Eng Appl Artif Intell 78:71–85. https://doi.org/10.1016/j.engappai.2018.10.014

Epp A (2016) Hashtag politics: a twitter sentiment analysis of the 2015 Canadian election using a randomized block design model

Gong Z, Cai T, Thill JC, Hale S, Graham M (2020) Measuring relative opinion from location-based social media: a case study of the 2016 us presidential election. Plos one 15(5):e0233660

Grčar M, Cherepnalkoski D, Mozetič I, Novak PK (2017) Stance and influence of twitter users regarding the Brexit referendum. Comput Soc Netw. https://doi.org/10.1186/s40649-017-0042-6

Gruhl D, Nagarajan M, Pieper J, Robson C, Sheth A (2010) Multimodal social intelligence in a real-time dashboard system. VLDB J 19(6):825–848

He J, Mao R, Shao Z, Zhu F (2020) Incremental learning in online scenario. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 13926–13935

He S, Zheng X, Zeng D, Cui K, Zhang Z, Luo C (2013) Identifying peer influence in online social networks using transfer entropy.

In: Pacific-Asia workshop on intelligence and security informatics, pp 47–61

Hussain M, Bird JJ, Faria DR (2018) A study on CNN transfer learning for image classification. In: UK workshop on computational intelligence, pp 191–202

Isotalo V, Saari P, Paasivaara M, Steineker A, Gloor PA (2016) Predicting 2016 us presidential election polls with online and media variables. In: Designing networks for innovation and improvisation. Springer, pp 45–53

Johnson R, Zhang T (2014) Effective use of word order for text categorization with convolutional neural networks. Preprint arXiv: 1412.1058

Kanoje S, Powar V, Mukhopadhyay D (2015) Using MongoDB for social networking website deciphering the pros and cons. In: 2015 international conference on innovations in information, embedded and communication systems (ICIIECS), pp 1–3

Kaplan AM, Haenlein M (2010) Users of the world, unite! the challenges and opportunities of social media. Bus Horizons 53(1):59–68

Karami A, Bennett LS, He X (2018) Mining public opinion about economic issues: Twitter and the us presidential election. Int J Strateg Decis Sci (IJSDS) 9(1):18–28

Katiyar A, Cardie C (2016) Investigating lstms for joint extraction of opinion entities and relations. In: Proceedings of the 54th annual meeting of the association for computational linguistics (Volume 1: Long Papers), pp 919–929

Krizhevsky A, Sutskever I, Hinton GE (2012) Imagenet classification with deep convolutional neural networks. In: Advances in neural information processing systems, pp 1097–1105

Kušen E, Strembeck M (2018) Politics, sentiments, and misinformation: an analysis of the Twitter discussion on the 2016 Austrian Presidential Elections. Online Soc Netw Media 5:37–50. https://doi.org/10.1016/j.osnem.2017.12.002

Lassen NB, la Cour L, Vatrapu R (2017) Predictive analytics with social media data. The SAGE handbook of social media research methods, p 328

Liu B (2012) Sentiment analysis and opinion mining. Synth Lect Hum Lang Technol 5(1):1–167

Liu B (2015) Sentiment analysis: mining opinions, sentiments, and emotions. Cambridge University Press

Liu R, Yao X, Guo C, Wei X (2020) Can we forecast presidential election using twitter data? an integrative modelling approach. Ann GIS:1–14

Liu W, Chen C, Wong KYK, Su Z, Han J (2016) STAR-net: a spatial attention residue network for scene text recognition. In: BMVC, vol 2, p 7

MacDonald R, Mao X (2015) Forecasting the 2015 general election with internet big data: an application of the trust framework

Muñoz-Expósito M, Oviedo-García MÁ, Castellanos-Verdugo M (2017) How to measure engagement in twitter: advancing a metric. Internet Research

Pang B, Lee L (2008) Opinion mining and sentiment analysis. Found Trends Inf Retriev 2(1–2):1–135

Paolanti M, Kaiser C, Schallner R, Frontoni E, Zingaretti P (2017) Visual and textual sentiment analysis of brand-related social media pictures using deep convolutional neural networks. In: International conference on image analysis and processing, pp 402–413

Park S, Ko M, Kim J, Liu Y, Song J (2011) The politics of comments: predicting political orientation of news stories with commenters' sentiment patterns. In: Proceedings of the ACM 2011 conference on computer supported cooperative work, pp 113–122.

Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/1958824.1958842

Park SJ, Lim YS, Sams S, Nam SM, Park HW (2011) Networked Politics on Cyworld: the text and sentiment of Korean political profiles. Soc Sci Comput Rev 29(3):288–299. https://doi.org/10.1177/0894439310382509

Peng S, Li J, Yang A (2015) Entropy-based social influence evaluation in mobile social networks. In: International conference on algorithms and architectures for parallel processing, pp 637–647

Peng S, Zhou Y, Cao L, Yu S, Niu J, Jia W (2018) Influence analysis in social networks: a survey. J Netw Comput Appl 106:17–32

Qian Q, Huang M, Lei J, Zhu X (2016) Linguistically regularized lstms for sentiment classification. Preprint arXiv:1611.03949

Rabade R, Mishra N, Sharma S (2014) Survey of influential user identification techniques in online social networks

Rai A, Borah S (2020) Study of various methods for tokenization

Santos CD, Gatti M (2014) Deep convolutional neural networks for sentiment analysis of short texts. In: Proceedings of COLING 2014, the 25th international conference on computational linguistics: technical papers, pp 69–78

Sapountzi A, Psannis KE (2018) Social networking data analysis tools & challenges. Fut Gen Comput Syst 86:893–913. https://doi.org/10.1016/j.future.2016.10.019

Schoen H, Gayo-Avello D, Metaxas PT, Mustafaraj E, Strohmaier M, Gloor P (2013) The power of prediction with social media. Internet Research

Shan G, Xu S, Yang L, Jia S, Xiang Y (2020) Learn#: a novel incremental learning method for text classification. Expert Syst Appl 147:113198

Simonyan K, Zisserman A (2014) Very deep convolutional networks for large-scale image recognition. Preprint arXiv:1409.1556

Stieglitz S, Dang-Xuan L (2013) Social media and political communication: a social media analytics framework. Soc Netw Anal Min 3(4):1277–1291. https://doi.org/10.1007/s13278-012-0079-3

Tsaytler P, Harding HP, Ron D, Bertolotti A (2011) Selective inhibition of a regulatory subunit of protein phosphatase 1 restores proteostasis. Science 332(6025):91–94

Wang C, Zhang P (2012) The evolution of social commerce: the people, management, technology, and information dimensions. Commun Assoc Inf Syst 31(1):5

Wu L, Barash G, Bartolini C (2007) A service-oriented architecture for business intelligence. In: IEEE international conference on service-oriented computing and applications (SOCA'07), pp 279–285

Yang M, Tu W, Wang J, Xu F, Chen X (2017) Attention-based LSTM for target-dependent sentiment classification. In: Proceedings of the thirty-first AAAI conference on artificial intelligence, pp 5013–5014

Yu S, Kak S (2012) A survey of prediction using social media. Preprint arXiv:1203.1647

Zafarani R, Abbasi MA, Liu H (2014) Social media mining: an introduction. Cambridge University Press

Zhang L, Wang S, Liu B (2018) Deep learning for sentiment analysis: a survey. Wiley Interdiscip Rev Data Min Knowl Discov 8(4):e1253