

Modelling Political Disaffection from Twitter Data

Corrado Monti
Dipartimento di Informatica,
Università degli Studi di Milano
me@corrado Monti.com

Alessandro Rozza
Dipartimento di Scienze
Applicate, Università degli
Studi di Napoli "Parthenope"
alessandro.rozza@
uniparthenope.it

Giovanni Zappella
Dipartimento di Matematica,
Università degli Studi di Milano
giovanni.zappella@unimi.it

Matteo Zignani
Dipartimento di Informatica,
Università degli Studi di Milano
matteo.zignani@unimi.it

Adam Arvidsson
Dipartimento di Scienze
Politiche, Università degli
Studi di Milano
adam.arvidsson@unimi.it

Elanor Colleoni
Department of Intercultural
Communication and
Management, Copenhagen
Business School
elc.ikl@cbs.dk

ABSTRACT

Twitter is one of the most popular micro-blogging services in the world, often studied in the context of political opinion mining for its peculiar nature of online public discussion platform. In our work we analyse the phenomenon of *political disaffection* defined as the "lack of confidence in the political process, politicians, and democratic institutions, but with no questioning of the political regime". Disaffection for organised political parties and institutions has been object of studies and media attention in several Western countries. Especially the Italian case has shown a wide diffusion of this attitude. For this reason, we collect a massive database of Italian Twitter data (about 35 millions of tweets from April 2012 to October 2012) and we exploit scalable state-of-the-art machine learning techniques to generate time-series concerning the political disaffection discourse.

In order to validate the quality of the time-series generated, we compare them with indicators of political disaffection from public opinion surveys. We find political disaffection on Twitter to be highly correlated with the indicators of political disaffection in the public opinion surveys. Moreover, we show the peaks in the time-series are often generated by external political events reported on the main newspapers.

General Terms

Political disaffection; Classification; Twitter; Sentiment Analysis

1. INTRODUCTION

Twitter is one of the most popular micro-blogging services in the world. Micro-blogging allows the publication of short text messages, used to share all kinds of information; on Twitter, these messages are called "tweets" (their maximum length is 140 characters), and many millions of them are posted every day.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WISDOM '13 August 12, 2013, Chicago, USA
Copyright 2013 ACM 978-1-4503-2332-1/13/08 ...\$15.00.

Twitter has proven to be a relevant data source to explore public sentiment trends ([4, 32]). Its content is easily available, and its flexible nature allows harvesting open conversations, public opinions, and news commentaries. Another crucial characteristic of Twitter is its timeliness; this peculiarity guarantees that tweets are related to a much narrower temporal window with respect to other user-generated texts, such as blogs.

Modelling trends from Twitter data has become a popular research task. Among such studies, those drawing attention to political topics are some of the most attractive, and in the last years a great deal of research works has focused on them.

In this study we concentrate on *political disaffection*, an important concept in political science. Political disaffection has been defined by Di Palma as "the subjective feeling of powerlessness, cynicism, and lack of confidence in the political process, politicians, and democratic institutions, but with no questioning of the political regime" [11]. In political science, levels of political disaffection are understood to relate to levels of political participation and, consequently they have important implications for the legitimacy of democratic political systems. This phenomenon has been significant since the 1960's in many Western countries and it has gained even more mass media attention after the 2008 Financial Crisis. These facts make the study of political disaffection pressing issue of contemporary studies of political behaviour. Some of the most relevant studies (i.e. [1, 25, 30]) are focused on Italy, since this political phenomenon has been particularly important in this country in the last half century. Moreover, the annual European Commission surveys on public opinion¹ confirm the relevance of this attitude among Italian citizens.

To our knowledge political disaffection has never been studied using Twitter data. In this work we propose an automatic approach to measure political disaffection using a massive collection of Twitter data from the Italian community. Our aim is the study of the relations between our measurement of political disaffection and political disaffection as measured by public opinion surveys.

In accordance with Di Palma we define political disaffection as negative sentiment towards the political system in general, rather than towards a particular politician, policy or issue. We operationalize this concept by defining expression of political disaffection tweets that have the following characteristics: 1) *political*, 2) *negative sentiment*, and 3) *generality*; where the last two features

¹http://ec.europa.eu/public_opinion

capture the lack of confidence in the whole political system. Consequently the measurement of political disaffection in Twitter can be performed by a sequence of three tasks. First, we use a supervised methodology to extract a subset of political tweets from the universe of tweets. Second, we perform a sentiment analysis on political tweets to extract those with negative sentiment. Third, we automatically select the tweets that refer to politics or politicians in general, rather than specific political events or personalities.

By applying our approach to the Italian Twitter community, we can monitor the trend of the political disaffection. This allows us to relate and compare the Twitter disaffection time series to indicators of political disaffection in public opinion surveys and thereby to validate our operational measurement. Finally, we show that external events such as important political news from Italian newspapers are often correlated with peaks in the produced time-series. This paper is organized as follows: in Section 2 the related works are summarized; in Section 3 we describe the procedures used to collect the datasets employed to train our supervised methods, the approach to extract the overall set of tweets employed in our analysis, and we summarize the public opinion surveys used to validate the quality of our approach; in Section 4 we describe the overall methodology to extract the political disaffection tweets; in Section 5 we present the achieved results on the extracted political disaffection trends; finally Section 6 contains the conclusions of our work.

2. RELATED WORKS

In literature a great deal of research has focused on the analysis of different phenomena using the data of micro-blogging services. Among them, in [24] the authors explore the correlation between types of user engagement and events about celebrities using Twitter data. Furthermore, in [3] the authors propose an approach to predict the stock market trend by monitoring micro-blogging.

The most closely related works are those concerning concept-level sentiment analysis [7], short text conceptualization [8] and the investigation of political topics using Twitter data. In [4], the authors propose a method to extract different time series corresponding to the evolution of 6 emotional attributes (tension, depression, anger, vigour, fatigue, and confusion) called Profile of Mood States (POMS). The authors apply POMS to suggest that socio-economic agitations caused significant fluctuations of the mood levels.

One of the earliest papers discussing the feasibility of using Twitter data as a substitute of traditional public opinion surveys is [22]. The authors employ Opinion-Finder² to determine both a positive and a negative score for each tweet in their dataset. Then, raw numbers of positive and negative tweets regarding a given topic are used to compute a sentiment score. Subsequently, sentiment time series are created for different topics such as: presidential approval, consumer confidence, and US 2008 Presidential elections. According to the authors both consumer confidence and presidential approval public opinion surveys show correlation with the Twitter sentiment data computed with their approach. However, no correlation has been found between electoral public opinion surveys and Twitter sentiment data.

In [31] an analysis of the tweets related to different parties running for the German 2009 Federal election has been carried out. The authors show that the volume of tweets mentioning a party or a candidate accurately reflected the election results, suggesting

²Opinion-Finder is a system that performs subjectivity analysis, automatically identifying when opinions, sentiments, speculations and other private states are present in text.

a possible approach to perform an electoral prediction. Furthermore, in [18] a novel method that aims at predicting elections has been proposed. This approach relies both on Twitter data and on additional information such as the party a candidate belongs to, or incumbency. Bermingham *et al.* [2] improve the previous approaches by incorporating sentiment analysis to the prediction of the political election. The authors tested their method in the 2011 Irish General Election finding that the results are not competitive when compared with traditional public opinion surveys. Similar approaches are proposed in [29, 28]. Nevertheless, the possibility to perform an electoral prediction using Twitter data [21, 13] is still an open issue. For instance, in [21] the authors analyse the results of different elections and they conclude that Twitter data is only slightly better than chance when predicting elections.

For this reason we avoid to predict election outcome in terms of percentages of party support, but we evaluate the well-known political attitude of political disaffection by analysing Twitter data through machine learning techniques. In order to validate the quality of the information extracted from the Twitter data, we highlight the relations of this data with political disaffection as measured in public opinion surveys.

3. DATA EMPLOYED

To model political disaffection on Twitter, we have to consider how a relevant tweet should be defined. According to the aforementioned Di Palma's definition, in the following we identified the characteristics of the tweets that express political disaffection:

- **Political:** obviously the subject of the tweets regards what Di Palma indicates as "political process, politicians and institution".
- **Negative:** we capture the powerlessness and the lack of confidence in the political system by analysing the negative feeling towards the subject of the tweets.
- **General:** the last sentence in Di Palma's definition denotes that the tweet subject is not a particular element of the political system. As a consequence tweets regarding most of parties or the whole political class are intended to be general, while tweets addressing a specific politician or institution do not belong to this category.

Once defined the subject of our work, we developed a filtering mechanism based on machine learning techniques with the purpose of extracting tweets from which a political disaffection feeling emerges. This filtering mechanism was constituted by a sequence of three tasks directly linked with the above characteristics.

The first task concerns the detection and the extraction of political tweets; in the second we retrieved negative tweets from the political ones, while in the last task we detected general tweets from those resulting after the second task (for details see Section 4).

In order to apply supervised machine learning algorithm to the first and the second task, we needed a reliable and big enough dataset to train our classifiers.

3.1 Training Twitter Data (TwitterTrainData)

We built the training set by a 2-step procedure involving a semi-automatic search method that employs the Twitter API v1 and a labeling phase guided by experts. The collection phase began at the beginning of April 2012 and ended at the beginning of June 2012. We collected about 120,000 of tweets and retweets. The selected tweets result from a geo-localized trending topic³ search and a tar-

³Trending topics are the most popular and talked-about words and phrases on Twitter for a specific time period.

	“positive”= +1	“negative”= -1
“political”= +1	7'965	4'544
“political”= -1	15'831	

Table 1: Label distribution of the TwitterTrainData after the majority voting.

geted search on political themes. In particular, at the end of each day we requested the top 10 trending topics of the Italian community. As most trending topics regard non-political arguments (i.e. celebrities, sports or viral hashtags), we selected the political content and a subset of the non-political. Furthermore, in order to have a more meaningful number of political tweets, we searched for tweets related to politicians, political news from Italian online newspapers and talk-shows. As query keywords we chose the Italian politicians’ surname, parties and organizations, the topics of the top news in the political section of online newspapers and the official hashtags of TV-talks. The resulting dataset consists of a large corpus of about 40,000 records each one composed by the tweet content, its date, and the keyword used in the search.

Once the dataset was collected, we started the labeling phase, employing the expertise of a pool of 40 Italian sociology and political science students. Each student was assigned a set of 3000 tweets to be classified by means of a web application. The annotators performed an intensive training phase where an extensive set of tweets representing the three typologies have been proposed to them. Three different labels have been associated to each tweet, the first is the *political* label (+1: political, -1: not political), the second label *sentiment*, coded only for political tweets, is about the feeling and can assume the values +1 (positive or neutral) or -1 (negative), while the third is the binary label *general* (+1: general, -1: specific)⁴.

The procedure was made so that each tweet was labelled by three different experts. This way we could increase the accuracy and the meaningfulness of the labelling process by removing tweets with disagreeing political labels. On the resulting set, we employed a majority voting approach for labelling the sentiment of the tweets. This choice was justified by the high level of agreement of the sentiment label. Indeed, by taking into account the Krippendorff’s alpha coefficient⁵, we obtained for the sentiment label $\alpha = 0.79$. In Table 1 we report the label distribution after applying the above procedure. Both political/not political classes and negative/positive classes within political tweets are quite balanced.

On the opposite, for the general label we obtain a low Krippendorff’s alpha coefficient (0.41). It can be argued that “general speech” is a concept that could not unequivocally and objectively be defined for human beings, for this reason we chose to discard this label, and we decided to simplify this task by defining a rule-based approach described in Section 4.

The final dataset (TwitterTrainData) is composed by 28,340 Italian labelled tweets on political and sentiment categories. To the best of our knowledge, it represents one of the biggest dataset containing tweets classified by experts.

⁴Two examples of Italian political tweets: “#ballaró: rappresentanti del nulla? Ci stiamo riprendendo i nostri diritti. State attenti, quello che avete visto è; solo l’inizio”. Political, general, negative; “#agora in 30 minuti: la donna ventriloquo, l’autolesionismo della sinistra rappresentata da #vendola e #fassino”. Political, specific, negative.

⁵Krippendorff’s alpha coefficient [17] is a statistical measure of the agreement achieved when many evaluators code the same set of units of analysis in terms of the values of a variable.

3.2 Training Newspaper Data (NewsTrainData)

The adoption of TwitterTrainData in the training phase could present some drawbacks due to the limited period it spans. For example, some important features to achieve a good classification considering a narrow retrieval period might lose their relevance in a wider period. These drawbacks result in a limited generalization power of the model employed to classify the political tweets. To improve the generality, we built up an additional dataset (NewsTrainData) containing all the article titles of different Italian newspapers (*Repubblica*, *il Manifesto*, and *Libero*) so that they span the whole spectrum of the political points of view from the Right to the Left wing. More precisely, using the feed RSS history, we selected all the articles from January 1st, 2012 to October 10th, 2012 extracting the news title, and we employed the categorization proposed by the newspaper to associate a label to the title. If a news belongs to the political category proposed by the newspaper we set the label to +1, otherwise -1. The resulting NewsTrainData is composed by 17,388 labelled newspaper titles, 10,670 of which political (61%).

3.3 Italian Twitter Community Data

To obtain general results on the political disaffection we performed our analysis on a large sample of the Italian Twitter community. To achieve this goal, we randomly extracted 50,000 Italian users, which posted at least one Italian tweet⁶ in a fixed temporal range (October 10th to October 30th). Moreover, to extend our sample we selected for each user all its Italian followers, thus producing a set of 261,313 users. Furthermore, we considered only the user profiles that have been created before April 4th (obtaining 167,557 users), to prevent the problem of the continuous growth of the Italian Twitter community, which could affect the quality of our political disaffection measure. Finally, we extracted all the tweets of each selected user, for the period of interest (April 4th, 2012 to October 10th, 2012), producing our final set composed by 35,882,423 tweets (TweetCorpus). According to an estimate⁷, it roughly represents more than 50% of the tweet volume posted in that period.

3.4 Public Opinion Surveys

Once detected all the tweets expressing political disaffection in TweetCorpus, we compared the resulting time series with the trend of some indicators extracted from public opinion surveys. The public opinion surveys have been collected by a global market research company (IPSOS) from April 11th, 2012 to October 10th, 2012. The sampling procedure consisted of a survey through CATI (computer-assisted telephone interview) of a representative sample of the Italian electorate. More precisely, almost every week, respondents were contacted with a quota sampling on fixed parameters (age, gender, education) using the technique of random digit dialing.

From public opinion surveys we extracted some indicators related to the phenomenon under investigation. An aspect of political disaffection that can be captured by public opinion survey is the attitude of political inefficacy. This attitude expresses the (subjective) sense of powerlessness of citizens in politics and the disbelief in the accountability of the political system and of all political parties. This definition led us to employ a measure of political inefficacy,

⁶To identify if a tweet is written in Italian we employ the Guess-Language library (<https://code.google.com/p/guess-language/>).

⁷<http://daily.wired.it/news/internet/2012/09/26/numeri-twitter-italia.html>

t_i	INEFFICACY	NO_VOTE
2012-04-11	14.86%	13.23%
2012-04-18	13.77%	14.53%
2012-05-02	22.20%	19.05%
2012-05-09	-	13.37%
2012-05-16	12.93%	13.34%
2012-05-23	16.31%	12.47%
2012-06-05	12.07%	13.31%
2012-06-06	11.03%	13.76%
2012-06-13	-	10.99%
2012-06-20	10.77%	13.08%
2012-06-26	6.91%	9.29%
2012-06-27	6.84%	13.09%
2012-07-04	-	11.88%
2012-07-11	7.87%	10.04%
2012-07-17	9.51%	13.64%
2012-07-18	6.00%	10.03%
2012-07-25	-	13.26%
2012-09-04	-	13.53%
2012-09-12	8.46%	11.22%
2012-09-19	-	12.75%
2012-09-25	9.44%	12.04%
2012-09-26	10.46%	12.74%
2012-10-03	-	12.87%
2012-10-10	11.76%	14.38%

Table 2: Public Opinion Surveys for INEFFICACY and NO_VOTE indicators.

that we called INEFFICACY, which represents the percentage of the respondents who expressed the minimum (equal to 1) propensity to vote (PTV, where the range is from 1-low to 10-high) for all political parties included in the surveys⁸. The PTV has been coded for each significative Italian party⁹. It is reasonable to assume that very low propensity to vote for any party is a reflection of the lack of accountability in the political system.

Together with political inefficacy we also used low intentions to vote as a proxy of political disaffection, which reflects citizens’ sense of powerlessness (see [30]). We captured the intention to not vote at the next election using the NO_VOTE indicator. This indicator includes the percentage of survey respondents that declare to have extremely low intention to vote at the next elections¹⁰ (see Table 2). In details, we considered the people that answered 1 (1 low - 10 high) to the question “How likely is it that you will vote at the next election?”.

We want to emphasize that these two indicators refer to different aspects of political disaffection. INEFFICACY measures the scepticism of a citizen towards the whole political class and capture an attitude, which might or might not be associated to a behaviour. On the other hand, NO_VOTE can capture the potential behavioural consequences of political disaffection. Note that the propensity of this last behaviour is influenced by a variety of factors, including the proximity of political elections.

⁸The total sample consists in 38, 537 respondents (~ 2267 respondents per poll).

⁹PD (Partito Democratico), PDL (Popolo della Libertà.), Lega Nord, IdV (Italia dei Valori), UDC (Unione di Centro), FLI (Futuro e Libertà.), SeL (Sinistra Ecologia Libertà.), and M5S (Movimento 5 Stelle).

¹⁰The total sample consists in 24, 971 respondents (~ 1040 respondents per poll).

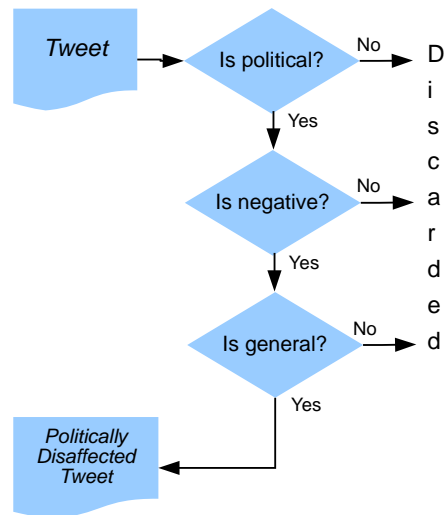


Figure 1: Classification chain employed.

4. CLASSIFICATION APPROACH FOR THE POLITICAL DISAFFECTION

Identifying political disaffection is a complex task even for human beings, so, in order to create a system for the detection of this attitude in tweets, we have to define it in a formal way. To that end, as described in Section 3, a political disaffection tweet has to match the following three criteria:

- **Political:** the tweet should regard politics.
- **Negative:** the sentiment of the tweet should be negative.
- **General:** the message have to regard politicians or parties in general. Tweets regarding only a political party or specific politician are not considered.

Since a classifier able to consider all this criteria at the same time can not be trained, we created a “chain” of classifiers as described in Figure 1. The relevant tweets after each step became the input for the next one. After every step the number of relevant tweets was less or equal to the number of relevant tweets after the previous step. The relevant tweets after the third step are eventually classified as *relevant* and all the other tweets are classified not relevant. Roughly speaking, we used TweetCorpus as input of the chain and we obtained a set of tweets denoting political disaffection (Tweet-Relevant) as output.

For the first step, we trained a classification algorithm using TwitterTrainData and NewsTrainData; the resulting classifier distinguished between political and non-political tweets. In the second step, the algorithm is trained with TwitterTrainData and the resulting classifier distinguishes between tweets with negative and non-negative sentiment. Please note that the TwitterTrainData collection is fixed, but the features are extracted in different ways depending on the classification step. The third and last step was performed by an ad-hoc classifier created with a rule-based approach to identify the general speech. Precisely, as noticed in Section 3, generality is a concept that could not unequivocally and objectively be defined for human beings, for this reason we opted to simplify this problem by employing a set of keywords identified by our experts to model

Keywords used for the general task	
politici	politicians
classe politica	political class
partiti	parties
deputati	members of parliament
senatori	members of senate
lo stato	the state
casta	clique

Table 3: The list of keywords selected by domain experts to model the general task in the political field.

this task (see Table 3). These keywords represent the most frequently used $\{1, 2\}$ -gram of words that refer to the political class in general. Furthermore, to improve the generalization of this approach we employed DBpedia. This database allows performing queries and provides a simple and automatic way to capture the semantic behind words based on Wikipedia. By using this database we extracted all the Italian politicians with their political affiliation (left, centre, and right) and we substituted their names with their affiliations. We selected as general those tweets that contained the keywords identified by experts or at the same time all the possible political affiliations (left, centre, and right)¹¹.

In the next sections we summarize the feature extraction methodologies, and, subsequently, the results of different classification approaches.

4.1 Feature Extraction Approaches

The efficacy of textual classification crucially depends on how the textual data is transformed into numerical features. Nevertheless, identifying the best method for feature extraction is a non-trivial problem, and the results are usually task dependent. For these reasons, we separately managed the two supervised classification tasks: political topics and negative sentiment. Note that in political topics we employed both the tweets data (TwitterTrainData) and newspaper titles (NewsTrainData).

We compared different techniques for features extraction in order to find the most suitable for our problems: n-grams of characters, single words¹², $\{1, 2, 3\}$ -grams of words, and we also applied more sophisticated approaches such string kernels [19].

For each of the techniques listed above, we computed: term frequency [20], boolean term presence [34], and term frequency-inverse document frequency (TF-IDF, [20]). Moreover, an important improvement was given by performing a stemming process and collapsing synonyms into a single feature. To perform this task we employed a freely available Italian synonyms dictionary¹³.

The test has been performed for the political topic classification using a 4-fold cross-validation with an online linear classifier¹⁴. Our results showed that 5-grams of characters constitutes our best option, independently of the counting scheme.

Taking into account the negative sentiment task and replicating the experiments with the same methodology of the previous case, we noted that the feature characterized by the space-separated word tokenization achieves the overall best results, employing as counting

¹¹The tweet: "#Bersani,#Berlusconi,#Monti sono tutti ladri" (#Bersani,#Berlusconi,#Monti are all thieves) becomes "left, centre, right are all thieves".

¹²Considering the space-separated words approach, we recognize as single word also emoticons and single punctuation marks such as "!". The URLs are also transformed in a unique token: (link).

¹³<http://webs.racocatala.cat/llengua/it/sinonimi.htm>

¹⁴Passive Aggressive, [10]




Figure 2: Linear fitting between the sentiment ratio (regressor) and the inefficacy indicator (response variable). The $R^2 = 0.61$ and the significance value for the sentiment ratio coefficient ($\rho = 3 \cdot 10^{-4}$) stress the goodness of the fitted model.

scheme the term frequency. In order to make our classifier more robust, we removed in each tweet the object the sentiment is referred to (sentiment target) extracting from DBpedia a full and up-to-date set of Italian political parties, politicians, and political offices. Combining this data with the recognition of strings starting with "@" (Twitter user-names) we were able to remove the sentiment target from the tweets.

4.2 Tested Classification Algorithms

To achieve our goal, we needed classifiers able to scale on huge corpus and possibly to be updated over time. Therefore, we especially focused our attention on online classifiers since they require only a single sweep on the data, making the classification process really fast with really good performances on the accuracy side.

We ran all the experiments on an ordinary workstation: Intel(R) Core(TM) i7-2600K CPU at 3.40GHz with 16Gb of RAM.

We tested four different online algorithms for classification¹⁵ and one batch classification algorithm:

- **ALMA** [14]: is a fast classifier which try to approximate the maximal margin hyperplane between the two classes. We set the parameter p equals to 2.
- **OIPCAC** [26]: is a classification method that employs a modified approach to estimate the Fisher Subspace, which allows to manage classification tasks where the space dimensionality is bigger than, or comparable to, the cardinality of the training set, and to deal with unbalanced classes.
- **PASSIVE AGGRESSIVE (PA, [10])**: is a Perceptron-like method. In our experiments we tested only the binary classifier with different settings.
- **PEGASOS** [27]: is a well-known online Support Vector Machine (SVM) solver.
- **RANDOM FOREST (RF, [6])**: is an algorithm based on an ensemble of classification trees. Since the algorithm is widely used in machine learning challenges with good results, we will use it as yardstick in our comparison.

¹⁵Most of these algorithms are well known and have a MATLAB implementation available in DOGMA [23].

Classifier	Accuracy	F-Measure	Global time
ALMA	0.883 ± 0.014	0.886 ± 0.011	13.5 ± 1
<i>PA</i>	<i>0.889 ± 0.012</i>	<i>0.890 ± 0.012</i>	<i>10.62 ± 0.1</i>
PEGASOS	0.882 ± 0.010	0.883 ± 0.010	1103 ± 10
OIPCAC	0.889 ± 0.001	0.891 ± 0.010	5911 ± 52

Table 4: 10-fold results for political topic detection (in bold face the best results considering F-measure). In italic the classifier selected for the classification process. With "time", we intend the time employed for training and classification, in seconds. We were not able to conclude all the runs with RF due to its high requirement of resources. Note that the number of samples is 45,728.

Classifier	Accuracy	F-Measure	Global time
<i>ALMA</i>	<i>0.703 ± 0.029</i>	<i>0.745 ± 0.034</i>	<i>0.82 ± 0.28</i>
PA	0.665 ± 0.064	0.705 ± 0.124	0.91 ± 10 ⁻³
PEGASOS	0.691 ± 0.033	0.732 ± 0.045	76 ± 0.1
OIPCAC	0.714 ± 0.026	0.751 ± 0.024	121 ± 25
RF	0.724 ± 0.026	0.776 ± 0.027	2173 ± 48

Table 5: 10-fold results for negative sentiment detection (in bold face the best results considering F-measure). In italic the classifier selected for the classification process. With "time", we intend the time employed for training and classification, in seconds. Note that the number of samples is 12,476.

We compared these algorithms on the two aforementioned tasks: political and negative. Note that we tested the online learning algorithms in a batch setting. In order to speed up the classification process we used only their one-sweep behaviour.

After an extensive tuning of the parameters, in Table 4 and in Table 5 we reported for each predictor its best performances in 10-fold cross validation on the political classification task and on the negative sentiment identification. In Table 4 the best result has been obtained by OIPCAC. The other classifiers achieved similar results, especially Passive Aggressive that further shows to be the fastest algorithm tested. In Table 5 the best result has been achieved by Random Forest, even though it had a very high running time. The variation between Table 4 and Table 5 highlighted that our sentiment classification task is more difficult w.r.t. the topic detection (political identification).

Considering the achieved results, in order to obtain a good trade off between accuracy and running times, we adopted the combination of PA and ALMA. In particular we used PA for the political/non-political classification and ALMA for the sentiment classification part. It is also possible to obtain comparable performances with different combinations of classifiers.

It is important to recall that since the general task (see Section 3) is difficult to be unequivocally modelled in all its aspect by a set of users, we simplified this problem by employing an approach based on a list of a few keywords selected by domain experts and by exploiting DBpedia as described above.

5. RESULTS

In this section we describe the time-series obtained employing the information extracted with the approach described in Section 4 and the relations between them and the public opinion surveys. Moreover, we summarize our methodology to identify the political news that produces the highest peaks of the generated time-series (breaking news).

To perform a correlation analysis with the INEFFICACY indicator taken from surveys, we employed the approach described in Section 4 to generate the set of tweets denoting political disaffec-

tion (TweetRelevant). Subsequently, taking into account each survey sampling date t_i (see Table 2), we generated three time-series computing the ratio between the number of political disaffection tweets and the number of political tweets by employing three time intervals:

1. from the date of the survey to 14 days before (Δ_1^{14});
2. from the day of the survey to 7 days before (Δ_1^7);
3. from 7 days before the date of the survey to 14 before (Δ_7^{14}).

The same approach has been employed for the NO_VOTE indicator.

Table 6 shows the Pearson correlation index computed between the political disaffection tweet-series and the INEFFICACY time-series. The best result (0.79) represents a strong correlation value between INEFFICACY and the information extracted by our approach. Furthermore, it is important to stress that the best time interval is Δ_1^{14} . We can strengthen our result by analysing the dependency between the political disaffection ratio and the INEFFICACY indicator. As shown in Figure 2, we find a linear dependence between the variables expressed by line $y = 16.6x + 0.01$, i.e. an increase of the political inefficacy corresponds to an increase in the Twitter political disaffection. These results combined with the Twitter timeliness suggest that our approach would be able to capture change in disaffection more promptly than public opinion surveys as can be noticed in Figure 3.

Table 7 shows the Pearson correlation index computed between the political disaffection tweet-series and the NO_VOTE one. The best result (0.59) represents a medium correlation value but still relevant showing that there is some connection between the modelled political disaffection and the intention to not participate at the next election day.

5.1 Breaking News Identification

After verifying the correlation between INEFFICACY and the Twitter political disaffection, we employed the TweetRelevant data to empirically determine some of the possible causes that produce the variation in disbelief in politics and politicians, hypothesising that citizens' political inefficacy is affected by controversial political news reported daily in the media. To achieve this goal we identified the peaks of the time-series generated as the daily ratio between the number of political disaffection tweets and the number of political tweets, and we associated each peak to a news belonging to NewsTrainData.

More precisely, to identify the peaks we employed an approach similar to that proposed in [15], taking into account as peaks the points of the time series greater than $\mu + 2\sigma$, where μ is the mean of the points of the time-series and σ is the standard deviation. However to improve the quality of our results we considered for each point a set of its neighbours¹⁶ (instead of all the points) to estimate the local μ s and σ s. The qualitative results are shown in Figure 4.

¹⁶We use a temporal window of 10 days, 5 before and 5 after the point of the time-series taken into account.

Interval	ρ	95% Confidence Interval	P-Value ($\rho > 0$)
Δ_1^{14}	0.7860	0.476-0.922	0.031%
Δ_7^{14}	0.7749	0.454-0.917	0.042%
Δ_1^7	0.6880	0.310-0.878	0.226%

Table 6: Pearson correlation index achieved between Twitter political disaffection and INEFFICACY time-series (p -value and confidence interval are two-tailed).




Figure 3: Twitter political disaffection time-series employing Δ_1^7 compared with the INEFFICACY indicator.




Figure 4: Political disaffection tweets day by day, with the selected peaks (highlighted by circles).

Interval	ρ	95% Confidence Interval	P-Value ($\rho > 0$)
Δ_1^{14}	0.5920	0.248-0.803	0.231%
Δ_7^{14}	0.5579	0.190-0.788	0.567%
Δ_1^7	0.4433	0.049-0.718	3.00%

Table 7: Pearson correlation index achieved between Twitter political disaffection and NO_VOTE time-series (p -value and confidence interval are two-tailed).

To associate each peak to a news, firstly we created an inverse document frequency (IDF) map by employing the words extracted from the corpus of the news included in NewsTrainData, and 1 million tweets randomly selected from the political subset (PTweetCorpus) of TweetCorpus (we employed the same classifier used for the political task described in Section 4 to identify the political tweets). Note that these weights reduce the relevance of the terms that are recurrent in many tweets. For each previously identified peak we create TD-IDF vectors for the tokenized news and tweets by employing the IDF map, thus obtaining, two vector sets for each day:

- \mathcal{N} the vectors' set of the news;
- \mathcal{T} the vectors' set of the tweets belonging to PTweetCorpus.

Subsequently we employed the cosine similarity between vectors to select the most correlated news with respect to the peak taken

into account as follows:

$$\arg \max_{n \in \mathcal{N}} \sum_{t \in \mathcal{T}} \frac{n \cdot t}{\|n\| \|t\|}$$

The results achieved are summarized in Table 8 where we reported the news with the highest cosine similarity. Table 8 evidences that disaffection peaks correlate to a broad and diverse range of political news. The news spans from the discussion about economical Italian crisis and the contested labor reform to bribe scandals and the individual behaviors of leaders or members of a specific political party.

Finally, we qualitatively compared the news identified with this approach with the trending topics on Twitter related to the day of each peak and we could note that most of the news effectively corresponded to one of the political daily trend. However, for few peaks, NewsTrainData did not contain any news correlated with the majority of the tweets of that day. Looking at the Twitter trending topics, it can be argued that this happened whenever the political discussions on Twitter did not concern any facts reported in newspapers, but the discussions spontaneously grew in the Twitter community. A meaningful example concerns the trending topic #no2giugno: this movement asked for the suspension of the military parade of June the 2nd (the Italian republic day), seen as a waste of resources, to use the money to rebuild the cities of Emilia (Italian region) after the earthquakes of 2012. This discussion generated two peaks (May 30th and 31st) that did not correlate with the traditional media news. A similar behaviour could be noted for the other two uncorrelated peaks.

2012-04-13	✓ La Lega prova a rifarsi un'immagine. Rinuncia agli ultimi rimborsi elettorali. Lega tries to clean up its image. It opts last electoral refunds out.
2012-04-17	✓ Lavoro, Monti pensa alla fiducia "I partiti approveranno la riforma". Labor, Monti thinks of a vote of confidence "Parties will enact reform".
2012-04-18	✓ Monti: niente crescita fino al 2013, disagio lavoro per metà famiglie. Monti: no economic growth until 2013, disadvantage for half of families.
2012-05-09	✓ Bersani: "Pd più forte, Monti ci ascolti" Grillo: "Partiti morti". Crollo del Pdl. Bersani: "PD stronger, Monti listen to us" Grillo: "Parties are dead". PDL falls.
2012-05-20	✓ Grillo su Brindisi: strage di Stato, fa comodo a loro. Grillo about Brindisi: state massacre, it's convenient for them.
2012-05-24	× Grillo attacca: "Noi soldi non li vogliamo rinunceremo a rimborsi prossime politiche". Grillo bashes: "We don't need money, we opt last electoral refunds out".
2012-05-30	× Riforma Csm, il gelo di Monti così è fallito il piano di Caticcalá. CSM reform, Monti's chill, Caticcalá's project is doomed.
2012-05-31	× Spread, Monti resta preoccupato "Rischio contagio malgrado gli sforzi". Spread, Monti worried "Risk contagion despite the efforts".
2012-07-14	✓ Cicchitto: "Primarie sono inutili Berlusconi candidato premier". Cicchitto: "Primary election is useless, Berlusconi is the premier candidate".
2012-07-19	× Monti ora teme il crac della Sicilia. Now Monti is afraid of Sicily default.
2012-08-09	✓ Monti al WSJ: "Con Berlusconi spread a 1200" L'ira del Pdl. E votano contro il governo. Monti to WSJ: "Spread at 1200 with Berlusconi" PDL anger and vote against government.
2012-08-28	✓ Grillo a Bersani: "Io fascista? Tu sei un fallito d'accordo con la P2". Grillo to Bersani: "Am I Fascist? You're failed at one with P2".
2012-09-23	✓ Così si rubava alla Regione Lazio ecco le rivelazioni di Fiorito ai pm. Fiorito's admission to public prosecutor: "How we stole moneys from Lazio government".
2012-09-26	✓ Caso Sallusti, salta l'accordo con il giudice il direttore domani rischia il carcere. Sallusti's instance, legal agreement breaks, the lead director risks the jail.
2012-09-27	✓ Sallusti: "In Italia mancano le palle" Paolo Berlusconi respinge le dimissioni. Sallusti: "In Italy many wimps" Paolo Berlusconi rejects Sallusti's resignation.

Table 8: Each row represents the news with the highest cosine similarity identified by our approach. The symbol ✓ represents the identified news that also appears in the political Twitter trending topic of the day taken into account. As additional contextual information, Lega is an Italian party, Bersani, Cicchitto and Berlusconi are politicians, PD is the Italian Democratic Party, Monti is the Italian ex-premier, Sallusti and Paolo Berlusconi (Silvio Berlusconi's brother) are respectively the lead director and the editor of a newspaper, Grillo is a comic/politician, Fiorito is a regional councilman involved in bribe inquiry. CSM is the magistrates' internal board of supervisors. P2 is a secret society.

6. CONCLUSIONS AND FUTURE WORKS

In this work we analysed the well-known attitude of political disaffection by using Twitter data through machine learning techniques. While the majority of the research has concentrated on the investigation of Twitter users' political attitude in term of support to a specific candidate or party [31, 2, 29, 28] to our knowledge, no prior studies have analysed manifestation of political disaffection. To investigate this phenomenon we concentrated on the Italian case, which has been identified in the political science literature as one of the most extreme cases of disaffection. We found evidence of a diffuse and consistent political discourse that can be classified as political disaffection. In order to validate the quality of the time-series generated employing our approach, we compared our result with political disaffection as measured in public opinion surveys (low intentions to vote and low political efficacy). We found evidence of a strong correlation between the two time-series. Fur-

thermore, we showed that important political news of Italian newspapers is often correlated with the highest peaks of the produced time-series. Taking into consideration the specificities of Twitter as a news medium (i.e. immediate diffusion of news and information), these results indicated that the data extracted from Twitter and from surveys are two reflections of a common underlying development that exhibits different temporal characteristics. This led us to suggest that Twitter data can be taken as a valid measurement of the fluctuating dimension of political disaffection.

This work showed that using Twitter to capture public opinion is more feasible than using it to predict the behavior of the public based on that opinion. Indeed, given the contradictory results of the electoral prediction based on Twitter data [21, 13], the different task of political disaffection seems to offer an interesting research topic for further investigations. This is suggested by the strong correlations between the time-series generated by employing the public opinion surveys and the Twitter data automatically extracted by our approach. Overall our results together with those presented in literature [4, 22] showed that for some phenomena the amount of Twitter discussion is a good measure of the diffusion of this phenomenon in society, despite the bias (such as age and geolocalization) of the Twitter population. The fact that we considered only people spontaneously writing on Twitter suggests that the resulting index is not a count of 'votes', but a measure of how much people are willing to spread these ideas in their lives (presumably not only using micro-blogs).

Moreover, Twitter's timeliness in relation to political events, with respect to traditional public opinion surveys, suggests that our method could be employed to perform a daily prediction of the citizen's political disaffection changes.

To further extend our approach and reach better results, we would like to improve our method to extract "generic speech" modelling this concept by means of ad-hoc ontologies [12]. Moreover we could enhance the classification accuracy with a proper selection of the tweets to be labelled by experts in an active-learning fashion (i.e. [9]), and we could improve the quality of the sentiment analysis by employing the top systems proposed in the SemEval 2013 competition. Furthermore, we could introduce the graph topology information in order to have a better understanding of the social component of this political phenomenon and the possibility to employ graph-based classifiers (i.e. [33], [16]).

Finally, since Twitter communications include a high percentage of ironic and sarcastic messages [28], another interesting improvement of our approach could be focused on the identification of these tweets as shown in [5].

7. REFERENCES

- [1] A. G. Almond and S. Verba. *The Civic Culture. Political Attitudes and Democracy in Five Nations*. SAGE Publications, Inc., 1963.
- [2] A. Bermingham and A. Smeaton. On using twitter to monitor political sentiment and predict election results. In *Proceedings of the Workshop on Sentiment Analysis where AI meets Psychology (SAAIP)*, 2011.
- [3] J. Bollen and H. Mao. Twitter mood as a stock market predictor. *Computer*, 44(10):91–94, 2011.
- [4] J. Bollen, A. Pepe, and H. Mao. Modeling public mood and emotion: Twitter sentiment and socio-economic phenomena. In *Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media (ICWSM)*, 2011.
- [5] C. Bosco, V. Patti, and A. Bolioli. Developing corpora for sentiment analysis and opinion mining: the case of irony and senti-tut. *Intelligent Systems, IEEE*, 28(2):53–63, 2013.

- [6] L. Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001.
- [7] E. Cambria, B. Schuller, Y. Xia, and C. Havasi. New avenues in opinion mining and sentiment analysis. *Intelligent Systems, IEEE*, 28(2):15–21, 2013.
- [8] E. Cambria, Y. Song, H. Wang, and N. Howard. Semantic multi-dimensional scaling for open-domain sentiment analysis. *Intelligent Systems, IEEE*, 2013.
- [9] N. Cesa-Bianchi, C. Gentile, and F. Orabona. Robust bounds for classification via selective sampling. In *Proceedings of the Annual International Conference on Machine Learning (ICML)*, 2009.
- [10] K. Crammer, O. Dekel, J. Keshet, S. Shalev-Shwartz, and Y. Singer. Online passive-aggressive algorithms. *Journal of Machine Learning Research*, 7:551–585, 2006.
- [11] G. Di Palma. Apathy and participation: Mass politics in western societies. *The Public Opinion Quarterly*, 36(2):297–299, 1972.
- [12] L. M. Garshol. Metadata? Thesauri? Taxonomies? Topic maps! Making sense of it all. *Journal of Information Science*, 30(4):378–391, 2004.
- [13] D. Gayo-Avello. I wanted to predict elections. *CoRR*, abs/1204.6441, 2012.
- [14] C. Gentile. A new approximate maximal margin classification algorithm. *Journal of Machine Learning Research*, 2:213–242, 2001.
- [15] D. Gruhl, R. Guha, D. Liben-Nowell, and A. Tomkins. Information diffusion through blogspace. In *Proceedings of the International conference on World Wide Web (WWW)*, 2004.
- [16] M. Herbster, M. Pontil, and S. R. Galeano. Fast prediction on a tree. In *Proceedings of the Conference on Neural Information Processing Systems (NIPS)*, 2008.
- [17] K. H. Krippendorff. *Content Analysis: An Introduction to Its Methodology*. Sage Publications, Inc, 2nd edition, 2003.
- [18] A. Livne, M. P. Simmons, E. Adar, and L. A. Adamic. The party is over here: Structure and content in the 2010 election. In *Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media (ICWSM)*, 2011.
- [19] H. Lodhi, C. Saunders, J. Shawe-Taylor, N. Cristianini, and C. J. C. H. Watkins. Text classification using string kernels. *Journal of Machine Learning Research*, 2:419–444, 2002.
- [20] C. D. Manning, P. Raghavan, and H. Schütze. *Introduction to information retrieval*. Cambridge University Press, 2008.
- [21] P. T. Metaxas, E. Mustafaraj, and D. Gayo-Avello. How (not) to predict elections. In *Proceedings of the Conference series on Social Computing (SOCIALCOM)*, 2011.
- [22] B. O’Connor, R. Balasubramanian, B. R. Routledge, and N. A. Smith. From Tweets to Polls: Linking Text Sentiment to Public Opinion Time Series. In *Proceedings of the International AAAI Conference on Weblogs and Social Media (ICWSM)*, 2010.
- [23] F. Orabona. *DOGMA: a MATLAB toolbox for Online Learning*, 2009. Software available at <http://dogma.sourceforge.net>.
- [24] A.-M. Popescu and M. Pennacchiotti. “dancing with the stars”, nba games, politics: An exploration of twitter users’ response to events. In *Proceedings of the International AAAI Conference on Weblogs and Social Media (ICWSM)*, 2011.
- [25] R. D. Putnam. *Making democracy work: civic traditions in modern Italy*. Princeton University Press, 1993.
- [26] A. Rozza, G. Lombardi, E. Casiraghi, and P. Campadelli. Novel fisher discriminant classifiers. *Pattern Recognition*, 45(10):3725–3737, 2012.
- [27] S. Shalev-Shwartz, Y. Singer, and N. Srebro. Pegasos: Primal estimated sub-gradient solver for svm. In *Proceedings of the Annual International Conference on Machine Learning (ICML)*, 2007.
- [28] M. Skoric, N. Poor, P. Achananuparp, E.-P. Lim, and J. Jiang. Tweets and votes: A study of the 2011 singapore general election. In *Proceedings of the 45th Hawaii International Conference on System Sciences (HICSS)*, 2012.
- [29] E. Tjong Kim Sang and J. Bos. Predicting the 2011 dutch senate election results with twitter. In *Proceedings of the Workshop on Semantic Analysis in Social Media (LASM)*, pages 53–60, Avignon, France, 2012.
- [30] M. Torcal and J. Montero. *Political Disaffection in Contemporary Democracies: Social Capital, Institutions and Politics*. Routledge, 2006.
- [31] A. Tumasjan, T. O. Sprenger, P. G. Sandner, and I. M. Welp. Predicting elections with twitter: What 140 characters reveal about political sentiment. In *Proceedings of the International AAAI Conference on Weblogs and Social Media (ICWSM)*, 2010.
- [32] N. Vallina-Rodriguez, S. Scellato, H. Haddadi, C. Forsell, J. Crowcroft, and C. Mascolo. Los twindignados: The rise of the indignados movement on twitter. In *Proceedings of the Conference series on Social Computing (SOCIALCOM)*, 2012.
- [33] F. Vitale, N. Cesa-Bianchi, C. Gentile, and G. Zappella. See the tree through the lines: The shazoo algorithm. In *Proceedings of the Conference on Neural Information Processing Systems (NIPS)*, 2011.
- [34] P. S. Yelena Mejova. Exploring feature definition and selection for sentiment classifiers. In *Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media (ICWSM)*, 2011.