Political Tendency Identification in Twitter using Sentiment Analysis Techniques

Ferran Pla and Lluís-F. Hurtado
Departament de Sistemes Informàtics i Computació
Universitat Politècnica de València
Camí Vera s/n 46022 València (Spain)
{lhurtado|fpla}@dsic.upv.es

Abstract

This paper describes an approach for political tendency identification of Twitter users. We define some metrics that take into account the polarity of the political entities in the tweets of each user. To obtain this polarities we present the sentiment analysis system developed. The evaluation was performed on the general corpus developed at TASS2013 workshop for Spanish. To our knowledge, the results obtained for the sentiment analysis task and the political tendency identification task are the best results published until now using this data set.

1 Introduction

Social media are usually used to express opinions and feelings about companies, products, services, hobbies, politics, etc. Therefore, enterprises, organizations, governments, and different groups in general have shown interest in the opinions that users have for their activities. They are also interested to known the way users use these media, the communication behaviour, and some users attributes such as gender, age, geographical location, political orientation, etc. In general, the main aim is to provide personalized services, particularized offers, or simply to know what people think about something in order to improve their activities.

The scientific community has made a great effort to provide effective solutions to analyse, structure, and process the large amount of on-line reviews in social media. A wide set of techniques of Sentiment Analysis (SA) are used in micro-blogging texts to extract the polarity (positive, negative, mixed or neutral) that users express in these texts. In this respect, Twitter has become a popular micro-blogging site in which users express their opinions on a variety of topics in real time. The texts used in Twitter are called tweets, which are short texts of a maximum of 140 characters and a language that does not have any restriction on the form and content. The nature of these texts poses new challenges for researchers in Natural Language Processing (NLP). In some cases, the tweets are written with ungrammatical sentences with a lot of emoticons, abbreviations, specific terminology, slang, etc. Therefore, the usual techniques of NLP must be adapted to these characteristics of the language, and new approaches must be proposed in order to successfully address this problem. NLP tools like POS taggers, parsers, or Named Entity Recognition (NER) tools usually fail when processing tweets because they generally are trained on grammatical texts and they perform poorly in micro-blogging texts.

In this work we present a system for addressing the task of political tendency identification of Twitter users based on SA techniques. For each user, we collect all their tweets and we extract all the entities related to the political subject. Then, we automatically assign a polarity to these entities and we define a political tendency metric that uses this entity polarity information combined with another tendency metric for classifying the political tendency of each user in four categories: *Left*, *Right*, *Center*, or *Undefined*. The evaluation of our system is performed on the General Corpus, a corpus of Spanish tweets provided by the organization of the TASS2013 workshop.

The paper is organized as follows. In section 2 we present relevant works for Twitter user classification and Sentiment Analysis. In Section 3 we present a description of the corpus used to evaluate our user

Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers, pages 183–192, Dublin, Ireland, August 23-29 2014.

political tendency system. This system is based on SA techniques. A description of our SA system is described in section 4. In Section 5 we describe the way we classify users according to their political leading. The evaluation and discussion of the results obtained are presented in section 6. Finally, in section 7 we present some conclusions and possible directions for future works.

2 Related works

The different approaches for estimating the political leaning of Twitter users explore features that range from *text content*, *users behavior* (taking into account the tweets and retweets information) and the *Twitter structure* (by considering the followers users, following users, etc.). An interesting study of some useful features to classify latent users attributes (gender, age, regional origin, and political orientation) is presented in (Rao et al., 2010). In (Conover et al., 2011a; Conover et al., 2011b) and study of the political alignment of Twitter users is performed by analyzing the way users communicates by means of retweets and user mentions. In (O'Connor et al., 2010a) SA techniques are used to determine the positive and negative polarity of Twitter messages. They also study the connexion between these polarities and the public opinion derived from traditional polling in order to substitute or complement them. (Pennacchiotti and Popescu, 2011) present a machine learning approach to Twitter user classification in democrats or republicans. With respect to the linguistic content they considered prototypical words and hashtags that are common in democrats or republicans users which provides clues for the classification. They also use SA tecniques based on lexicons for the classification task. In (Boutet et al., 2012) polical leading of users is performed by counting the amount of tweets related to political parties analysing the hashtags. They also consider the interaction among parties by analyzing the retweets and mentions. Users interaction by analysing tweets and retweets is also the main idea of the work presented in (Wong et al., 2013).

In (Cohen and Ruths, 2013) previous works on political orientation of Twitter users are analyzed to conclude that the accuracy results reported are overstimated do to the way the data sets are constructed. When these approaches are applied to normal Twitter users accuracy results significantly decrease.

Sentiment Analysis (SA) has been widely studied in the last decade in multiple domains. Most work focuses on classifying the polarity of the texts as positive, negative, mixed, or neutral. The pioneering works in this field used supervised (Pang et al., 2002) or unsupervised (knowledge-based) (Turney, 2002) approaches. In (Pang et al., 2002), the performance of different classifiers on movie reviews was evaluated. In (Turney, 2002), some patterns containing POS information were used to identify subjective sentences in reviews to then estimate their semantic orientation.

The construction of polarity lexicons is another widely explored field of research. Opinion lexicons have been obtained for English language (Liu et al., 2005) (Wilson et al., 2005) and also for Spanish language (Perez-Rosas et al., 2012). A good presentation of the SA problem and a description of the state-of-the-art of the more relevant approaches to SA can be found in (Liu, 2012). An overview of the current state of different approaches to the subjectivity and SA task is presented in (Montoyo et al., 2012).

Research works about SA on Twitter are much more recent. Twitter appeared in the year 2006 and the early works in this field are from 2009 when Twitter started to achieve popularity. Some of the most significant works are (Barbosa and Feng, 2010), (Jansen et al., 2009), and (O'Connor et al., 2010b). A survey of the most relevant approaches to SA on Twitter can be see in (Martínez-Cámara et al., 2012) , (Vinodhini and Chandrasekaran, 2012). The SemEval2013 competition has also dedicated a specific task for SA on Twitter (Wilson et al., 2013), which shows the great interest of the scientific community in this field. The TASS2013 workshop has proposed different tasks for SA and political tendency identification focused on the Spanish language (Villena-Román and García-Morera, 2013).

3 The Corpus

The General Corpus of TASS2013[1](Villena-Román and García-Morera, 2013) contains approximately 68000 Twitter messages (*tweets*) written in Spanish (between November 2011 and March 2012) by 158 well-known personalities of the world of politics, economy, communication, mass media, and culture.

[1]This corpus is freely available on the web page of TASS2013 (http://www.daedalus.es/TASS2013).

The corpus is encoded in XML. Each tweet includes its ID (*tweetid*), the creation date (*date*), and the user ID (*user*). It is tagged with its global polarity using N and N+ labels for negative polarity with different intensity, P and P+ labels for positive polarity with different intensity, and the NEU label for neutral polarity. Label NONE was used to represent tweets with no polarity at all. Moreover, the polarity to the entities that are mentioned in the tweet was also included. The level of agreement of the expressed sentiment is annotated both for global and entity level. Also, a selection of a set of topics was made based on the thematic areas covered by the corpus, such as politics, soccer, literature, entertainment, etc. Each message is also assigned to one or several of these topics.

	N	N+	NEU	NONE	P	P+
training	1,335 (18.49%)	847 (11.73%)	670 (9.28%)	1,483 (20.54%)	1,232 (17.07%)	1,652 (22.88%)
test	11,287 (18.56%)	4,557 (7.50%)	1,305 (2.15%)	21,416 (35.22%)	1,488 (2.45%)	20,745 (34.12%)

Table 1: The distribution of the polarity of the tweets in the corpus.

Table 1 shows the distribution of tweets per polarity in the corpus. It is divided into two sets: training (about 10%, 7219 tweets) and test (about 90%, 60798 tweets). It can be observed that this distribution is not balanced for the different polarities. Finally, each user from the test set of the General corpus is labeled with their political tendency in four possible values: $Left$, $Right$, $Centre$, and $Undefined$.

4 Description and Evaluation of the Sentiment Analysis System

Figure 1 shows an overview of our system for the SA problem. The system consists of 4 modules. The first module is the Pre-processing module, which performs the tokenization, lemmatization, and Named Entities recognition of the input tweet. A lemma reduction and a POS tagging process is also carried out in this module. The second module is optional. It allows us to obtain the polarity of the entities contained in the tweet. If we omitted this step the global polarity of the tweet is obtained. The third module is the Feature Extraction module, which selects the features from the pre-processed tweet (or from the segments of tweets) and obtains a feature vector. Some features require the use of a polarity lexicon of lemmas and words. To determine the best features, a tuning process is required during the training phase. The fourth module is the Polarity Classifier module, which uses a classifier (learned from feature vectors of the training set) to assign a polarity label to the tweet.

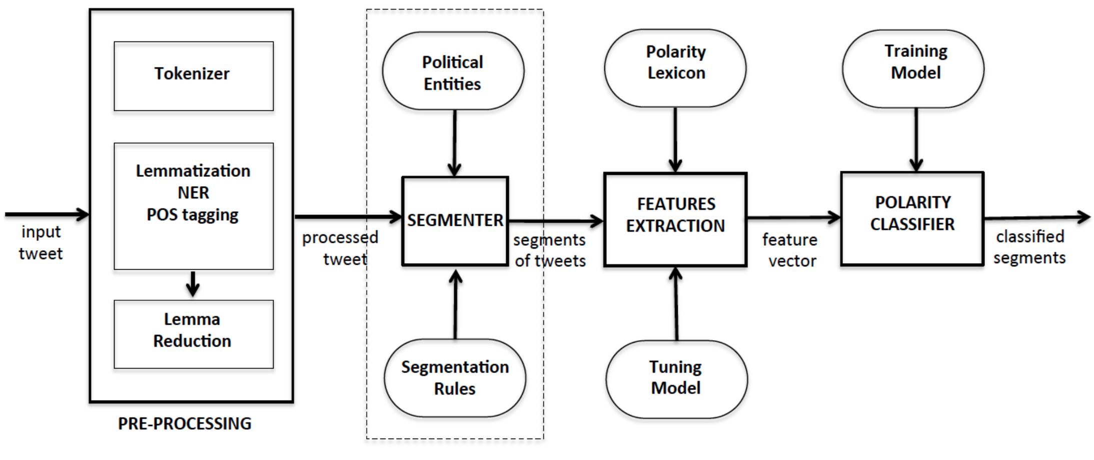

Figure 1: Sentiment Analysis System Overview

4.1 Pre-processing of Tweets

Before addressing the SA task, it is necessary to make a proper tokenization of the tweets. Although there are a lot of tokenizers available on the web, they need to be adapted in order to address the segmentation of tokens of a tweet. Furthermore, most of these resources are for the English language, which adds a degree of difficulty for their use in processing Spanish tweets.

Moreover, the use of NLP resources such as stemmers, POS taggers, parsers, NER systems, Word Sense Disambiguation (WSD) systems, etc. are impractical if the characteristics of the tweets are not taken into account. Therefore, an adjustment and adaptation must be made for the Twitter domain.

In our system, we decided to use and adapt available tools for tokenization, lemmatization, NER, and POS tagging. We adapted the package *Tweetmotif* [2] that is described in (O'Connor et al., 2010b) to process Spanish tweets. We also used *Freeling* [3] (Padró and Stanilovsky, 2012) (with the appropriate modifications for handling Twitter messages) for stemming, Named Entity Recognition, and POS tagging.

We added some functions to process special tokens (e.g., grouping all *hashtags* into a single token, grouping all *web* addresses into a single token or grouping all *url* into a single token). We also grouped the *dates* into a single token, the *numbers* into a single token, and the *punctuation* marks into a single token.

4.2 The Segmenter

For the proposed approach we need to determine the polarity of political entities that contains a tweet. It is because the polarity of each entity could be different of the global polarity of the tweet. In the tweet of the corpus[4]: *"Rajoy's government goes up the pensions. PSOE cuts back all things except the unemployment."* we have two entities, *Rajoy* (the president of Spanish government from the right-wing party PP) and *PSOE* (a Spanish left-wing party). This tweet is labeled with a neutral global polarity, but each entity have a different polarity (ENTITY (Rajoy): *Positive*. ENTITY (PSOE): *Negative*).

Even for tweets with only one entity we must decide what fragments of text refers to that entity. In the example: *"Rajoy already has been talking for an hour. Not saving anywhere only expenses, all reforms cost a lot of money. Did he tell us something at the end?"*, to determine the polarity of entity Rajoy, we must take into account all the tweet, because the two last sentences references to *ENTITY(Rajoy)*. In contrast, in the example: *"Today 349 members attending to the formation of the lower house. Only the AMAIUR deputy for Navarra is missing"*, only the sentence containing the AMAIUR entity is being required to determine its polarity.

Obtaining the polarity at entity level is a hard problem and introduces additional complexity because the part of the tweet refers to each of the entities must be determined. To resolve this problem it should make a deep parsing of the tweet and perform a study of such dependencies. This is not a solved problem in NLP even considering normative texts and is further aggravated in Twitter texts. Besides, in many cases, the dependencies are between different sentences, and problems such as coreference must be taking into account in order to determine, for example, which pronoun refers to a certain entity. Other problems such as synonyms and acronyms of certain entities can make this problem harder.

We have chosen a more simple and practice approach that consists in defining a set of heuristics to determine which segment of the tweet refers to each of the entities present on it. We defined some rules to do this segmentation. If the tweet contains only one entity the context considered was all the tweet. We evaluated other alternatives, but due to the short length of tweets, with that decision the best global results were obtained. If the tweet contains two entities, the casuistry is greater. If both entities are placed together at the beginning or the end of the tweet all the tweet is considered as a context for both entities. By contrast, if separate, and has sufficient context, the tweet is segmented by defining the context of each entity. Next, we show some examples and the segmentation produced by the defined rules.

Example 1 is the easier case due the two entities are in separated sentences. When both entities are in the same sentence, in Example 2 the rule applied determines that the context for the first entity is from the beginning until the second entity, and the rest of the sentence is the context for the second entity. Example 3 is more difficult, and the rules applied produce segmentations like this *[On March 25 we elect between the immobility of the @PSOE] [and the renovation and the hope of the @ppandaluz.])*, that are not correct but can be useful for determining the polarity of each entity. In addition, due to the short length of the tweets, the context of an entity is often so small that it does not contain information enough

[2]https://github.com/brendano/tweetmotif.

[3]http://nlp.lsi.upc.edu/freeling/

[4]All the exemples have been translated to English.

Example 1
*[***Rajoy***'s government goes up the pensions.] [***PSOE*** cuts back all things except the unemployment.]*
GLOBAL POLARITY: *NEU*. ENTITY (Rajoy): *Positive*. ENTITY (PSOE): *Negative*
Example 2
[As **IU** *gains confidence in Andalucía] [***PP*** loses members.]*
GLOBAL POLARITY: *NEU*. ENTITY (IU): *Positive*. ENTITY (PP): *Negative*
Example 3
On March 25 we elect between [the immobility of the **@PSOE***] and [the renovation and the hope of the* **@ppandaluz.]**
GLOBAL POLARITY: *NEU*. ENTITY (@PSOE): *Negative*. ENTITY (@ppandaluz): *Positive*

to correctly classify the polarity of the entity. In such case, the option that was chosen is to establish a threshold of context, and if it is below than this threshold, it was assigned the same polarity to all the entities of the tweet. When the number of entities is greater than two in much cases we assigned the same polarity to all the entities of the tweet because we had not enough context.

4.3 Feature Selection

The feature selection process was performed by cross validation (10-fold validation) using the training set to select the set of relevant features.

We considered the following set of features: unigrams and bigrams of lemmas obtained in the preprocessing of the tweets that belong to a set of selected POS. We considered only the lemmas of a minimum frequency (*f*) in the training set. We unified all *hashtags*, *user references*, *dates*, *punctuations* as a single feature. We classified the emoticons in the following categories: *happy, sad, tongue, wink, and other*. Finally, we used external polarity lexicons of lemmas and words.

Some of the features required further adjustment. For the POS feature we selected the lemas that belongs to the *nouns, verbs, adjectives, and adverbs* POS and also *exclamations* and *emoticons*. We estimated the minimum frequency of the lemmas to be selected (*f =2*). Finally, we selected the external lexicons to be used. One of the lexicons used was originally for English language (Wilson et al., 2005) that was translated into Spanish automatically, and other (Perez-Rosas et al., 2012) lexicon was a list of words that was originally in Spanish. Then, we combined these two resource with the lexicon presented in (Saralegi and San Vicente, 2013).

4.4 Polarity Classifier

The task was addressed as a classification problem that consisted of determining the polarity of each tweet. We used WEKA[5], which is a tool that includes (among other utilities) a collection of machine-learning algorithms that can be used for classification tasks. Specifically, we used a SVM-based approach because it is a well-founded formalism, that has been successfully used in many classification problems. In the SA task, SVM has shown it ability to handle large feature spaces and to determine the relevant features (Joachims, 1998).

We used the NU-SVM algorithm (Schölkopf et al., 2000) from an external library called *LibSVM*[6], which is very efficient software for building SVM classifiers. It is easy to integrate this software with WEKA thus allowing us to use all of WEKA's features. We used the *bag_of_words* approach to represent each tweet as a feature vector that contains the frequency of the selected features of the training set.

4.5 Evaluation of the Sentiment Analysis System

We evaluated our system on the SA tasks defined at the TASS2013 workshop. Two different sub-tasks called *5-level* and *3-level* were proposed. Both sub-tasks differ only in the polarity granularity considered. The *5-level* sub-task uses the labels *N*, *N+*, *P*, *P+*, and *NEU*. The *3-level* sub-task uses the labels *N*, *P*, and *NEU*. In both sub-tasks, an additional label (*NONE*) was used to represent tweets with no polarity.

The accuracy results obtained on the unseen data test were: 62.88%±0.38% for *5-level* task and 70.25%±0.36% for *3-level task*. This results outperformed all the approaches at TASS2013 workshop with statistical significance (with a 95% level of confidence). The official results ranged from 61.6% to 13.5% for the *5-level* task and from 66.3% to 38.8% for the *3-level* task. The F_1 result obtained in the

[5]http://www.cs.waikato.ac.nz/ml/weka/
[6]http://www.cs.iastate.edu/~ yasser/wlsvm/

Sentiment Analysis at Entity level task was worse (F_1=0.40), but it still is the best result reported in the sentiment analysis at entity level task at TASS2013 competition.

5 Political tendency identification

The objective of this task is to estimate the political tendency of each user from the test set of the General corpus in four possible values: $Left$, $Right$, $Centre$, and $Undefined$. Next, we describe the approach we proposed for this task. This approach uses the SA system previously described in section 4.

To perform the classification of users we assume the following hypothesis: the positive opinions on a political party is a political orientation similarly to the user performing the review for this party, on the contrary, a negative opinion about a party is a political orientation opposite to that shown by this party.

In this way, to classify users by their political orientation, first we identify entities associated with political parties and secondly we analyze the polarity of these entities in the tweets of each user.

We consider three types of entities: entities labeled by *Freeling* as proper names (i.e., comité_del_pp_de_madrid), Twitter users (i.e., @38congresopsoe), and Twitter hashtags (i.e., #upyd). Among all possible entities we selected those containing the acronym for a political party or the name of a political leader. A total of 864 entities related to political parties and political leaders were detected. Table 2 shows the parties and political leaders considered and some examples of the selected entities.

Party	**Tendency**	**Examples of Entities**
PP	right	#17congesoPP, congreso_nacional_pp, ppopular, congresopp, #ppfachas
PSOE	left	elpsoe, #adiosalpsoeenandalucia, #38congresopsoe
IU	left	asamblea_de_iu, iumalaga, diputados_de_iu, #iu
UPyD	centre	upydeuskadi, #demagogiaupyd, #mareamagenta, upyd_asturias
CiU	right	ciu+tripartito, #ciu, ciu-mintiendo-crujen

Political Leader	**Party**	**Examples of Entities**
Rajoy	PP	#rajoynoeslasolución, españa_de_rajoy, irpf_de_rajoy
González Pons	PP	@gonzalezpons, rajoy_para_gonzález_pons
Rubalcaba	PSOE	#rubalcabaenlaser, @conrubalcaba, rubalcaba_para_el_psoe
Zapatero	PSOE	nueva_via_de_zapatero, presidente_zapatero, zapatero_tv
Cayo Lara	IU	@cayo_lara, cayo_lara, cayo

Table 2: Tendency of political parties and political leaders.

We defined a tendency measure $Tendency$ that assigned a value of -1 to those entities related to left parties, a value of $+1$ to entities related to right parties and a value of 0 to the entities related to centre parties.

Next we show how has been numerically calculated the political orientation of users. For each user U_i of the General corpus we obtain the set T_i that includes all of their tweets that contain political entities. For users who do not have any tweet that contain political entities the *Undefined* label is assigned.

For each tweet $T_{i_j} \in T_i$, $j = 1 \cdots |T_i|$, we identify the political entities that are contained on it. Let E_{i_j} be the set of entities of the tweet T_{i_j}. We denote each of the entities contained in E_{i_j} as $E_{i_{j_k}} \in E_{i_j}, k = 1 \cdots |E_{i_j}|$.

We obtained the polarity of each entity by using the system described in section 4. After that, we assigned a numerical value to each polarity. In this respect, we assigned $Polarity = +1$ to the entities with positive polarity (label P), $Polarity = -1$ to the entities with negative polarity (label N) and finally, $Polarity = 0$ to the entities without polarity, that is, to the NEU and NONE labels.

We combined[7] the $Tendency$ and $Polarity$ measures previously presented to define a new measure ($Political\_Tendency$) to obtain the political orientation of each user.

[7]We have considered multiple combination strategies, in this work we present the combination with the best results.

$$Political\_Tendency(U_i) = \frac{\sum_{j=1\cdots|T_i|} \sum_{k=1\cdots|E_{i_j}|} Polarity(E_{i_{j_k}}) \cdot Tendency(E_{i_{j_k}})}{\sum_{j=1\cdots|T_i|} |E_{i_j}|} \quad (1)$$

From the $Political\_Tendency$ values obtained for each user, we classified the user tendency tanking into account the following: users without political entities in their tweets are classified as $Undefined$; users with $Political\_Tendency$ between -0.05 and +0.05 are classified as $Centre$; users with $Political\_Tendency$ lower than -0.05 are classified as $Left$; and users with $Political\_Tendency$ greater than +0.05 are classified as $Right$.

6 Experimental Evaluation of the Political Identification System

The measures selected to evaluate our approach were the Precision, the Recall, and the F-measure for $\beta = 1$ (F_1). Table 3 summarizes the experimental results of our proposal. The table includes both the overall results (*Global*) and the results for each one of the political tendencies (*Left*, *Right*, *Centre*, and *Undefined*). It also includes the distribution of the tendencies in the gold-standard (%Ref). For the global result, the precision and the recall are the same since each user in the test set had a tendency assigned and the task consist to assign a tendency to all the users.

Tendency	%Ref	Precision	Recall	F_1
Left	21.5	0.658	0.735	0.694
Centre	17.7	0.478	0.393	0.431
Right	39.9	0.786	0.698	0.739
Undefined	20.9	0.780	0.970	0.865
Global	100	0.709	0.709	0.709

Table 3: Experimental results obtained in the political tendency identification task of TASS2013.

The result obtained by our system (0.709) is the best result reported so far for this corpus, to our knowledge. The tendency for what we get better results is the *Undefined* (F_1=0.865). We consider the political tendency of a user to be *Undefined* if he did not have any tweet that references any of the majority parties. This assumption may be too strict for common users, but it seems reasonable for the well-know users that form the test corpus.

The tendency that our system had more trouble identifying was *Centre* (F_1=0.431). The tendency of a user can be identified as *Centre* when he expressed -in his tweets- opinions about entities related to centre parties, even when these opinions were negative. This is because the neutral value of *Centre* entities. In addition, users with opinions on right and left parties with the same polarity may be identified as *Centre*, which can be wrong in many cases.

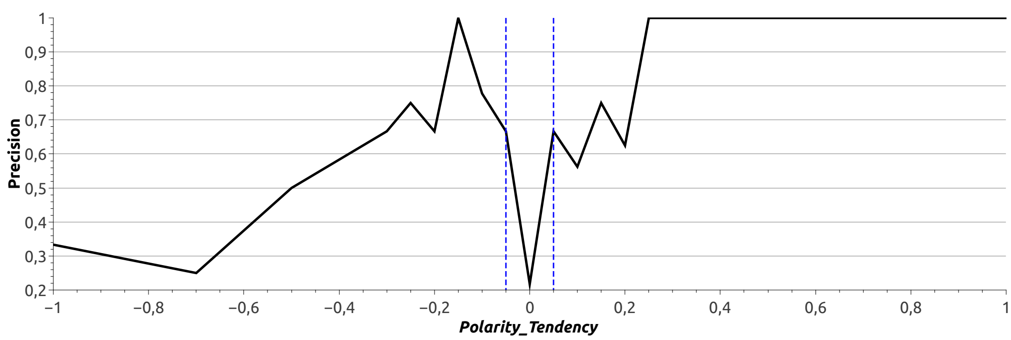

Figure 2: Precision results depending on the $Political\_Tendency$ assigned by the system.

Although it seems that the ability of our system to identify *Left* and *Right* tendencies was similar (F_1=0.694 for *Left* and F_1=0.739 for *Right*), analyzing the results considering the values of *Political_Tendency* some significant differences can be observed. Figure 2 shows the results, in terms of Precision, considering the value of *Political_Tendency* assigned to each user by our system, from a value of -1 (the maximum value for *Left*) to a value of +1 (the maximum for *Right*).

As expected, most identification errors occurred for *Political_Tendency* value near zero, should remember that values between -0.05 and 0.05 were considered *Centre*. Considering the *Right* tendency, all users that obtained *Political_Tendency* value greater than 0.25 were correctly identified as *Right*, performed better than would be expected. However, the behavior of the *Left* tendency was not symmetrical. It seems that values between -0.3 and -0.1 were better to determine correctly this tendency.

Although we have no clear explanation for this behavior, it could be due to multiple factors, including: the simplicity of the proposal, labeling errors in the polarity of certain entities, or the greater difficulty of numerically identify the *Left* tendency (at least in this corpus).

7 Conclusions

We have described our approach for political tendency identification of Twitter users. We have defined a metric, called *Political_Tendency*, that takes into account the polarity of entities related to political parties that appear in the tweets of the user. The Sentiment Analysis system developed in order to obtain the polarity of these entities was also presented.

The evaluation was performed using a corpus of Spanish tweets developed at TASS2013 workshop. This corpus was used for a specific political tendency identification task at this workshop. To our knowledge, the results obtained by our system are the best results published until now using this corpus.

We are very interested in SA tasks and in identifying tendencies in social media. In this sense, we have several ideas on how to improve our approach to identifying the political tendency in Twitter.

It would be interesting to test our approach using a larger corpus of tweets from normal user. We think that the characteristics of the users of the test corpus -figures of culture, journalism and politics in Spain- made the task a little easier. Perhaps the political tendency of ordinary users would be more difficult to identify. Moreover, the political spectrum would be more diverse and should increase the catalog of political parties. Moreover, the political spectrum would be more varied and, consequently, the catalog of political parties should be increased.

It should be emphasized the difficulty of building an annotated corpus of tweets that could be used to evaluate and compare different alternative systems. A great effort of acquisition of the tweets and a subsequent manual labeling process is required. In addition, a validation process is needed to correct the errors introduced by manual labeling. Even using crowdsourcing-based solutions it is a very expensive task both in money and time. In this context, to have a labeled corpus as the one provided by TASS2013 is a great help for the scientific community.

On the portability of the system, we think that it will be easy to adapt our proposal to another political context. This adaptation should focus on two different aspects. First, the Sentiment Analysis System should be adapted to a new language. In the case of languages with linguistic resources freely available the adaptation would be very simple. Second, political entities should be changed to fit the political context where we want to test the system. It would be sufficient to identify the most relevant parties and their leaders and classify them according to their political tendency. However, it is possible that in other political contexts different to Spanish, the *Left*, *Centre*, and *Right* tendencies also need to be adapted.

Finally, we have interest in using Machine Learning techniques for the task of identifying political tendency on twitter. On this point, we are working on a system in which *Political_Tendency*, as defined in this paper, will be a feature within a wider classification system. In this new system, we want to include additional information (not available in the TASS2013 corpus) about user behavior and Twitter structure in order to improve our approach.

Acknowledgements

This work has been partially funded by the Spanish MEC projects DIANA (TIN2012-38603-C02-01) and Tímpano (TIN2011-28169-C05-01).

References

Luciano Barbosa and Junlan Feng. 2010. Robust sentiment detection on twitter from biased and noisy data. In *Proceedings of the 23rd International Conference on Computational Linguistics: Posters*, pages 36–44. Association for Computational Linguistics.

Antoine Boutet, Hyoungshick Kim, and Eiko Yoneki. 2012. What's in Your Tweets? I Know Who You Supported in the UK 2010 General Election. In *The International AAAI Conference on Weblogs and Social Media (ICWSM)*, Dublin, Irlande, June.

Raviv Cohen and Derek Ruths. 2013. Classifying political orientation on twitter: It's not easy! In *International AAAI Conference on Weblogs and Social Media*.

M Conover, B Gonçalves, J Ratkiewicz, A Flammini, and F Menczer. 2011a. Predicting the political alignment of twitter users. In *Proceedings of 3rd IEEE Conference on Social Computing (SocialCom)*.

Michael Conover, Jacob Ratkiewicz, Matthew Francisco, Bruno Gonçalves, Alessandro Flammini, and Filippo Menczer. 2011b. Political polarization on twitter. In *Proc. 5th International AAAI Conference on Weblogs and Social Media (ICWSM)*.

Bernard J Jansen, Mimi Zhang, Kate Sobel, and Abdur Chowdury. 2009. Twitter power: Tweets as electronic word of mouth. *Journal of the American society for information science and technology*, 60(11):2169–2188.

Thorsten Joachims. 1998. Text categorization with support vector machines: learning with many relevant features. In Claire Nédellec and Céline Rouveirol, editors, *Proceedings of ECML-98, 10th European Conference on Machine Learning*, number 1398, pages 137–142, Chemnitz, DE. Springer Verlag, Heidelberg, DE.

Bing Liu, Minqing Hu, and Junsheng Cheng. 2005. Opinion observer: Analyzing and comparing opinions on the web. In *Proceedings of the 14th International Conference on World Wide Web*, WWW '05, pages 342–351, New York, NY, USA. ACM.

Bing Liu. 2012. *Sentiment Analysis and Opinion Mining. A Comprehensive Introduction and Survey*. Morgan & Claypool Publishers.

Eugenio Martínez-Cámara, M. Teresa Martín-Valdivia, L. Alfonso Ureña-López, and Arturo Montejo-Raéz. 2012. Sentiment analysis in twitter. *Natural Language Engineering*, 1(1):1–28.

Andrés Montoyo, Patricio Martínez-Barco, and Alexandra Balahur. 2012. Subjectivity and sentiment analysis: An overview of the current state of the area and envisaged developments. *Decision Support Systems*, 53(4):675–679.

Brendan O'Connor, Ramnath Balasubramanyan, Bryan R. Routledge, and Noah A. Smith. 2010a. From tweets to polls: Linking text sentiment to public opinion time series. In William W. Cohen and Samuel Gosling, editors, *ICWSM*. The AAAI Press.

Brendan O'Connor, Michel Krieger, and David Ahn. 2010b. Tweetmotif: Exploratory search and topic summarization for twitter. In William W. Cohen and Samuel Gosling, editors, *Proceedings of the Fourth International Conference on Weblogs and Social Media, ICWSM 2010, Washington, DC, USA, May 23-26, 2010*. The AAAI Press.

Lluís Padró and Evgeny Stanilovsky. 2012. Freeling 3.0: Towards wider multilinguality. In *Proceedings of the Language Resources and Evaluation Conference (LREC 2012)*, Istanbul, Turkey, May. ELRA.

Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. 2002. Thumbs up? sentiment classification using machine learning techniques. In *IN PROCEEDINGS OF EMNLP*, pages 79–86.

Marco Pennacchiotti and Ana-Maria Popescu. 2011. A machine learning approach to twitter user classification. In Lada A. Adamic, Ricardo A. Baeza-Yates, and Scott Counts, editors, *ICWSM*. The AAAI Press.

Veronica Perez-Rosas, Carmen Banea, and Rada Mihalcea. 2012. Learning sentiment lexicons in spanish. In Nicoletta Calzolari (Conference Chair), Khalid Choukri, Thierry Declerck, Mehmet Uğur Doğan, Bente Maegaard, Joseph Mariani, Jan Odijk, and Stelios Piperidis, editors, *Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC'12)*, Istanbul, Turkey, may. European Language Resources Association (ELRA).

Delip Rao, David Yarowsky, Abhishek Shreevats, and Manaswi Gupta. 2010. Classifying latent user attributes in twitter. In *Proceedings of the 2Nd International Workshop on Search and Mining User-generated Contents*, SMUC '10, pages 37–44, New York, NY, USA. ACM.

Xabier Saralegi and Iñaki San Vicente. 2013. Elhuyar at tass 2013. In *Proceedings of the TASS workshop at SEPLN 2013*. IV Congreso Español de Informática.

Bernhard Schölkopf, Alex J. Smola, Robert C. Williamson, and Peter L. Bartlett. 2000. New support vector algorithms. *Neural Comput.*, 12(5):1207–1245, May.

Peter D. Turney. 2002. Thumbs up or thumbs down? semantic orientation applied to unsupervised classification of reviews. In *ACL*, pages 417–424.

Julio Villena-Román and Janine García-Morera. 2013. Workshop on sentiment analysis at sepln 2013: An over view. In *Proceedings of the TASS workshop at SEPLN 2013*. IV Congreso Español de Informática.

G Vinodhini and RM Chandrasekaran. 2012. Sentiment analysis and opinion mining: A survey. *International Journal*, 2(6).

Theresa Wilson, Paul Hoffmann, Swapna Somasundaran, Jason Kessler, Janyce Wiebe, Yejin Choi, Claire Cardie, Ellen Riloff, and Siddharth Patwardhan. 2005. Opinionfinder: A system for subjectivity analysis. In *Proceedings of HLT/EMNLP on Interactive Demonstrations*, pages 34–35. Association for Computational Linguistics.

Theresa Wilson, Zornitsa Kozareva, Preslav Nakov, Sara Rosenthal, Veselin Stoyanov, and Alan Ritter. 2013. Semeval-2013 task 2: Sentiment analysis in twitter. *Proceedings of the International Workshop on Semantic Evaluation, SemEval*, 13.

Felix Ming Fai Wong, Chee Wei Tan, Soumya Sen, and Mung Chiang. 2013. Quantifying political leaning from tweets and retweets. In Emre Kiciman, Nicole B. Ellison, Bernie Hogan, Paul Resnick, and Ian Soboroff, editors, *ICWSM*. The AAAI Press.