10.12871/clicit2014218

Relying on intrinsic word features to characterise subjectivity, polarity and irony of Tweets*

Francesco Barbieri, Francesco Ronzano, Horacio Saggion
Universitat Pompeu Fabra, Barcelona, Spain
name.surname@upf.edu

Abstract

English. We describe our participation to the SENTIPOLC task of EVALITA 2014. We experimented the use of intrinsic word features to characterise each Tweet. We relied only on these features to train a set of Decision Trees to characterise the subjectivity, the polarity and the ironic contents of each Tweet. In Task 1 and Task 2 our model shows good performances comparing to the other participants, even if there is still space for improvements. In Task 3 our model do not achieve acceptable performances. We interpret and discuss these results.

Italiano. *Descriviamo la nostra partecipazione a SENTIPOLC di EVALITA 2014. Abbiamo sperimentato l'uso di features intrinseche delle parole per caratterizzare ogni Tweet. Grazie a queste features abbiamo costruito Decision Trees per determinare la soggettività, la polarità e il contenuto ironico di ogni Tweet. Nel Task 1 e Task 2 il nostro modello mostra buone prestazioni rispetto agli altri partecipanti, anche se c'è ancora spazio per miglioramenti. Nel Task 3 il nostro modello non raggiungere prestazioni accettabili. Nel paper discutiamo tali risultati fornendo possibili interpretazioni.*

1 Motivation

The automatic identification of the diverse facets of sentiments and opinions expressed by social media users constitutes a relevant and challenging research trend. In this context, the Sentiment Polarity Classification task of EVALITA 2014 (SENTIPOLC, Basile et al. (2014)) offers both a shared dataset and a venue to experiment and compare new approaches to the analysis of opinionated texts in social media. SENTIPOLC proposes three tasks respectively devoted to automatically determine the subjectivity, the polarity and the irony of a Tweet. This paper describes our participation in these three SENTIPOLC taks. We exploited an extended version of the Tweet classification features and approach described in Barbieri et al. (2014). In particular, we experimented the use of intrinsic word features, characterising each word in a Tweet (like usage frequency in a reference corpus, number of associated synsets, etc.), to try to model and thus automatically determine its subjectivity, polarity and ironic traits. We did not exploit textual features (like word occurrences, bigrams, skipgrams or other word patterns) to try to reduce the dependency of our model on a specific topic or on the set of words used in the considered domain. We aim to detect two aspects of Tweets by intrinsic word features: the style used (e.g. register used, frequent or rare words, positive or negative words, etc.) and the unexpectedness in the use of words, particularly important for subjectivity and irony (Lucariello, 1994). We exploited Decision Trees to classify Tweets in all the three SENTIPOLC tasks. In Section 2 we describe the tools we used to process Tweet contents. In Section 3 we introduce the features we built our model on. Section 4 analyses the performances of our model concerning the three tasks of SENTIPOLC.

2 Text Analysis and Tools

In order to process the text of Tweets so as to enable the feature extraction process, we used a set of freely available tools. First of all, we associated to each Tweet a normalised version of its text by expanding abbreviations and slang expressions, deleting emoticons, properly converting hashtags

*The research described in this paper is partially funded by the Spanish fellowship RYC-2009-04291, the SKATER-TALN_UPF project (TIN2012-38584-C06-03), and the EU project Dr. Inventor (n. 611383).

into words whether they have a syntactic role. We then tokenised, PartOfSpeech-tagged, applied Word Sense Disambiguation (UKB) and removed stop words from the normalized text of Tweets by exploiting Freeling (Carreras et al., 2004). We also used the Italian WordNet 1.6[1] to get synsets and synonyms of each word of a Tweet as well as the sentiment lexicon Sentix[2] (Basile and Nissim, 2013) derived from SentiWordnet (Esuli and Sebastiani, 2006) to get the polarity of synsets. We relied on the CoLFIS Corpus of Written Italian[3] to obtain the usage frequency of words in written Italian. We exploited the results of these analyses of the contents of Tweets to generate the word intrinsic features we describe in Section 3.

3 Our Model

In the three tasks of SENTIPOLC, we trained a Decision Tree to classify Tweets as far as concern their subjectivity, polarity and ironic contents. We exploited the widespread machine learning framework Weka in order to train and test our classification models. We characterised each Tweet by six classes of features all describing intrinsic aspects of the words of the same Tweet. These feature classes are: Frequency, Synonyms, Ambiguity, Part of Speech, Sentiments, and Punctuation.

3.1 Frequency

We accessed the CoLFIS Corpus to retrieve the frequency of each word of a Tweet. Thus, we derive three types of Frequency features: *rarest word frequency* (frequency of the most rare word included in the Tweet), *frequency mean* (the arithmetic average of all the frequency of the words in the Tweet) and *frequency gap* (the difference between the two previous features). These features are computed including all the words of each Tweet. We also determined these features by considering only Nouns, Verbs, Adjectives, and Adverbs.

3.2 Synonyms

We consider the frequencies (in CoLFIS Corpus) of the synonyms of each word in the Tweet, as retrieved from the Italian WordNet 1.6. Then we computed, across all the words of the Tweet: the *greatest / lowest number of synonyms* with frequency higher than the one present in the Tweet, the *mean number of synonyms* with frequency greater / lower than the frequency of the related word present in the Tweet. We determine also the greatest / lowest number of synonyms and the mean number of synonyms of the words with frequency greater / lower than the one present in the the Tweet (*gap* feature). We computed the set of Synonyms features by considering both all words of the Tweet together and only words belonging to each one of the four Parts of Speech listed before.

3.3 Ambiguity

To model the ambiguity of the words in the Tweets we use the WordNet synsets associated to each word. Our hypothesis is that if a word has many meanings (synset associated) it is more likely to be used in an ambiguous way. For each Tweet we calculate the *maximum number of synsets* associated to a single word, the *mean synset number* of all the words, and the *synset gap* that is the difference between the two previous features. We determine the value of these features by including all the words of a Tweet as well as by considering only Nouns, Verbs, Adjectives or Adverbs.

3.4 Part Of Speech

The features included in the Part Of Speech (POS) group are designed to capture the style of the Tweets. The features of this group are eight and each one of them counts the number of occurrences of words characterised by a certain POS. The eight POS considered are *Verbs*, *Nouns*, *Adjectives*, *Adverbs*, *Interjections*, *Determiners*, *Pronouns*, and *Appositions*.

3.5 Sentiments

The sentiments of the words in Tweets are important for two reasons: to detect the *sentiment* style (e.g. if Tweets contain mainly positive or negative terms) and to capture unexpectedness created by a negative word in a positive context and viceversa. Relying on Sentix (see Section 2) we computed the *number of positive / negative words*, the *sum of the intensities of the positive / negative scores of words*, the *mean of positive / negative score of words*, the *greatest positive / negative score*, the *gap between the greatest positive / negative score and the positive / negative mean*. As previously done, we computed these features by considering only Nouns, Verbs, Adjetives, and Adverbs.

[1] http://multiwordnet.fbk.eu/english/home.php
[2] http://www.let.rug.nl/basile/twita/sentix.php
[3] http://linguistica.sns.it/CoLFIS/Home_eng.htm

3.6 Punctuation

We also want to capture the punctuation style of the authors of a Tweet. Punctuation is very important in social networks: a full stop at the end of a subjective message may change the polarity of the message, the use of ellipses can be sign of irony (Carvalho et al., 2009). Each feature that is part of the Punctuation set is the number of a specific punctuation mark, including: ".", "#", "!", "?", "$", "%", "&", "+", "-", "=", "/".

	P	**R**	**F1**
Task 1 (subj.)	0.7332	0.6011	0.6497
Task 2 (polarity)	0.6565	0.5723	0.6049
Task 3 (irony)	0.5797	0.4591	0.4987

Table 1: Final scores (arithmetic average of the score of each class) of the three tasks organised in Precision, Recall and F-Measure.

4 Experiments and Results

In this section we present our results in the three SENTIPOLC tasks (see Table 1). We only report final results (mean of Precision, Recall and F-Measure of each class). In order to get other participants results, please refer to the SENTIPOLC paper (Basile et al., 2014).

4.1 Task 1: Subjectivity Classification

Given a message, decide whether the message is subjective or objective. Our model scores at position four out of nine groups. Our score is six points less than the best one in F-measure. Our system showed that we can determine if a Tweet is subjective or not with an acceptable precision by not considering explicitly words or word patterns, but only relying on intrinsic word features.

4.2 Task 2: Polarity Classification

Given a message, decide whether the message is of positive, negative, neutral or mixed sentiment (i.e. conveying both a positive and negative sentiment). In this task our model ranks fifth out of eleven participants. We obtain an averaged F-measure of 0.6049.

4.3 Task 3: Irony Detection

Given a message, decide whether the message is ironic or not. At this task our system scored as the last one, clearly showing that, at least for the Tweet dataset exploited in SENTIPOLC, relying only on intrinsic word features has limited power in determining if a Tweet is ironic or not.

5 Conclusions

In this paper we describe our participation to the SENTIPOLC task of EVALITA 2014. We experimented the use of intrinsic word features to characterise each Tweet. We relied exclusively on these features to train a set of Decision Trees respectively useful to determine the subjectivity, polarity and irony in Tweets. We explicitly decided not to rely on explicit words or word patterns as features. In Task 1 and Task 2 our model shows good performances comparing to other models, even if there is still space for improvements. In Task 3 our model do not achieve acceptable performances. Among other considerations, we related this issue to the fact that the training data in SENTIPOLC are strongly dependent on a specific topic, politics and this topic dependence limits the effectiveness of our system. In fact our classifier does not use words or word patterns that usually constitute features exploited to characterise a domain. In general, we noticed that avoiding text features may constitute a limitation for a classifier if the dataset to deal with concerns a specific topic and thus topic specific words could constitute good features to model the domain. As future work we are planning to experiment with other classification approaches (Support Vector Machines among them) as well as to evaluate the utility to complement the feature set we presented in this paper with word and word pattern features (like word occurrences, bigrams, skip-grams or other word patterns).

References

Francesco Barbieri, Francesco Ronzano, and Horacio Saggion. 2014. Italian Irony Detection in Twitter, a First Approach. In *Proceedings of the First Italian Conference of Computational Linguistic*, Pisa, Italy, December.

Valerio Basile and Malvina Nissim. 2013. Sentiment analysis on italian tweets. In *Proceedings of the 4th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 100–107.

Valerio Basile, Andrea Bolioli, Malvina Nissim, Viviana Patti, and Paolo Rosso. 2014. Overview of the Evalita 2014 SENTIment POLarity Classification Task. In *Proceedings of the 4th evaluation cam-*

paign of Natural Language Processing and Speech tools for Italian (EVALITA'14), Pisa, Italy.

Xavier Carreras, Isaac Chao, Lluis Padró, and Muntsa Padró. 2004. Freeling: An open-source suite of language analyzers. In *LREC*.

Paula Carvalho, Luís Sarmento, Mário J Silva, and Eugénio de Oliveira. 2009. Clues for detecting irony in user-generated contents: oh...!! it's so easy;-). In *Proceedings of the 1st international CIKM workshop on Topic-sentiment analysis for mass opinion*, pages 53–56. ACM.

Andrea Esuli and Fabrizio Sebastiani. 2006. Sentiwordnet: A publicly available lexical resource for opinion mining. In *Proceedings of Language Resources and Evaluation Conference*, volume 6, pages 417–422.

Joan Lucariello. 1994. Situational irony: A concept of events gone awry. *Journal of Experimental Psychology: General*, 123(2):129.