SSA-UO: Unsupervised Twitter Sentiment Analysis

Reynier Ortega, Adrian Fonseca
CERPAMID, University of Oriente
Ave Patricio Lumumba S/N
Santiago de Cuba, Cuba

Yoan Gutiérrez
DI, University of Matanzas
Autopista a Varadero Km $3\frac{1}{2}$
Matanzas, Cuba

Andrés Montoyo
DLSI, University of Alicante
Carretera de San Vicente S/N
Alicante,Spain

Abstract

This paper describes the specifications and results of SSA-UO, unsupervised system, presented in SemEval 2013 for Sentiment Analysis in Twitter (Task 2) (Wilson et al., 2013). The proposal system includes three phases: data preprocessing, contextual word polarity detection and message classification. The preprocessing phase comprises treatment of emoticon, slang terms, lemmatization and POS-tagging. Word polarity detection is carried out taking into account the sentiment associated with the context in which it appears. For this, we use a new contextual sentiment classification method based on coarse-grained word sense disambiguation, using WordNet (Miller, 1995) and a coarse-grained sense inventory (sentiment inventory) built up from SentiWordNet (Baccianella et al., 2010). Finally, the overall sentiment is determined using a rule-based classifier. As it may be observed, the results obtained for Twitter and SMS sentiment classification are good considering that our proposal is unsupervised.

1 Introduction

The explosion of Web 2.0 has marked a new age for the human society. The huge use of Social Media such as Facebook[1] , MySpace[2] , LinkedIn[3] and Twitter[4] , offers a place for people to share information in real time. Twitter is one of the most popular social network websites and has been growing at a very fast pace. The number of active users exceeds 500 million and the number of tweets posted by day exceeds 500 million (as of May 2012)[5]. Through the twitter applications, users shared opinions about personalities, politicians, products, companies, events, etc. This has been attracting the attention of different research communities interested in analyzing its content and motivated many natural language tasks, such as sentiment analysis, emotions detection, opinions retrieval, product recommendation or opinion summarization.

One of the most popular sentiment analysis tasks is polarity classification. This task is a new field that classifies opinion texts as positive, negative or neutral (Pang et al., 2002; Turney, 2002; Esuli and Sebastiani, 2006; Wilson et al., 2006; Wiegand et al., 2010). Determining polarity might seem an easy task, as many words have some polarity by themselves. However, words do not always express the same sentiment, and in most cases the polarity of a word depends on the context in which the word is used. So, terms that clearly denote negative feelings can be neutral, or even positive, depending on their context. Hence, sentiment analysis systems should include semantic-level analysis in order to solve word ambiguity and correctly capture the meaning of each word according to its context. Also, complex linguistic processing is needed to deal with problems such as the effect of negations and informal language. Moreover, understanding the sentimental meaning of the different textual units is important to accurately determine the overall polarity

[1]https://www.facebook.com
[2]http://www.myspace.com/
[3]http://www.linkedin.com
[4]https://www.twitter.com/
[5]http://www.statisticbrain.com/twitter-statistics/

*Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 2: Seventh International Workshop on Semantic Evaluation (SemEval 2013)*, pages 501–507, Atlanta, Georgia, June 14-15, 2013.

of a text.

In this paper, we present a system that has as main objective to analyze the sentiments of tweets and classify these as positive, negative or neutral. The proposal system includes three phases: data preprocessing, contextual word polarity detection and message classification. The preprocessing phase comprises treatment of emoticons, spell-errors, slang terms, lemmatization and POS-tagging. Word polarity detection is carried out taking into account the sentiment associated with the context within which it appears. For this, we use a new contextual sentiment classification method based on coarse-grained word sense disambiguation, using WordNet (Miller, 1995) and a coarse-grained sense inventory (sentiment inventory) built up from SentiWordNet (Baccianella et al., 2010). Finally, the polarity is determined using a rule-based classifier. The paper is organized as follows. Section 2 describes of SSA-UO system. In Section 3 we evaluate our proposal and discuss the results obtained in the SemEval 2013 Task No. 2. Finally, section 4 provides concluding remarks.

2 SSA-UO System

We use an unsupervised strategy consisting in a coarse-grained clustering-based word sense disambiguation (WSD) method that differentiates positive, negative, highly positive, highly negative and objective uses of every word on context which it occurs. The proposal method uses WordNet and a coarse-grained sense inventory (sentiment inventory) built up from SentiWordNet. The overall architecture of our sentiment classifier is shown in Figure 1.

Firstly, data preprocessing is done to eliminate incomplete, noisy or inconsistent information. A Sentiment Word Sense Disambiguation method (Section 2.3) is then applied to content words (nouns, adjectives, verbs and adverbs). Once all content words are disambiguated, we apply a rule-based classifier (Section 2.4) to decide whether the tweet is positive, negative or neutral.

Unsupervised word sense disambiguation method proposed by (Anaya-Sánchez et al., 2006) was adapted for sentiment word sense disambiguation. Unlike the authors, who aim to obtain the correct sense of a word, we use the method to determine

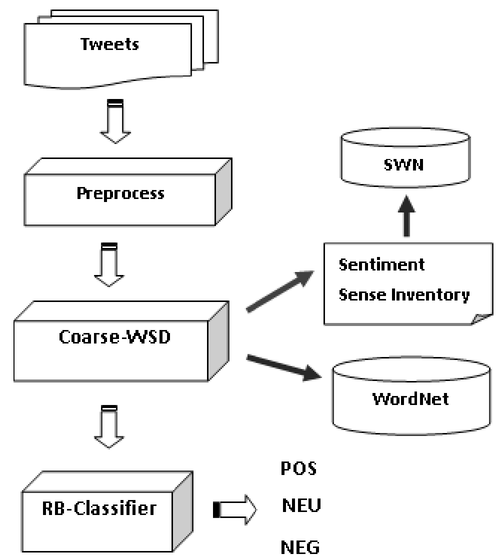

Figure 1: Overall architecture of Sentiment Classifier

when a word is used with highly positive (HP), positive (P), highly negative (HN), negative (N) or objective (O) meaning based on a sentiment sense inventory. We make sentiment sense inventory based on sense-level annotation in SentiWordNet. Finally, we apply a rule-based classifier to determine the overall sentiment in tweet.

2.1 Data Preprocessing

The tweets differ from the text in articles, books, or even spoken language. It is limited to 140 characters, also includes many idiosyncratic uses, such as emoticons, slang terms, misspellings, URLs, “RT” for re-tweet, “@” for user mentions, “#” for hashtags, and character repetitions. Therefore it is necessary to preprocess the text, in order to reduce the noise information. The preprocessing step involve the following task. The text is tokenized and URL, re-tweets and author mentions are removed. Hashtag tokens frequently contain relevant information related to the topic of the tweet, this is included as part of the text but without the “#” character. We replace emoticon tokens by emotion words using an emoticons dictionary, obtained from Wikipedia

[6]. Each emoticon was manually annotated with an emotion word and polarity value. Emoticons that suggest positive emotions - ":-)", ":)", "X-D" - are annotated with the emotion word *"happy"* and negative emoticons - ":-(", ":-c", ":,(" - are annotated with the emotion word *"sad"*. The presence of abbreviations within a tweet is noted, therefore abbreviations are replaced by their meaning (e.g., LOL – laughing out loud) using a dictionary[7]. Finally the text is POS-tagged and lemmatized using TreeTagger (Schmid, 1994) and stopwords are discarded.

2.2 Sentiment Sense Inventory

We considered SentiWordNet for building sentiment coarse-grained sense inventory. SentiWordNet contain positive, negative and objective scores between 0 and 1 for all senses in WordNet. Based on this sense level annotation, we define a new rule (*SentiS*) for classifying senses in five sentiment class. The senses are classified in the following manner (Alexandra et al., 2009): senses whose positive score is greater than or equal to 0.75 are considered to be highly positive (HP), senses with positive score greater than or equal to 0.5 and lower than 0.75 are considered positive (P), senses with negative score greater than or equal 0.75 are considered highly negative (HN), whereas those whose negative score is lower than 0.75 and greater than or equal to 0.5 are considered to be negative (N). In the remaining cases, the senses are considered to be objective (O) (see equation(1)).

$$sentiS(s) = \begin{cases} HP\ if\ ScoreP(s) \geq 0.75 \\ HN\ if\ ScoreN(s) \geq 0.75 \\ P\ if\ ScoreP(s) < 0.75\ and\ ScoreP(s) \geq 0.5 \\ N\ if\ ScoreN(s) < 0.75\ and\ ScoreN(s) \geq 0.5 \\ O\ in\ other\ case \end{cases} \tag{1}$$

Table 1 summarizes the distribution of the five sentiment classes once classified all senses of SentiWordNet.

A notable unbalance can be observed between the number of highly positive, highly negative, positive, negative and objective senses.

[6]http://en.wikipedia.org/wiki/List_of_emoticons

[7]http://www.noslang.com/dictionary/

Once all senses were classified in a five sentiment sense class, we create a coarse sense inventory based on this classification. This inventory is defined in the following manner: For each word in SentiWordNet we grouped its senses with the same sentiment class in a single sense (coarse-sense), in case of objective senses these are kept separated.

2.3 Contextual Word Polarity Detection

Much work on sentiment analysis have been directed to determine the polarity of opinion using anotated lexicons with prior polarity (Hatzivassiloglou and McKeown, 1997; Kamps and Marx, 2002; Turney, 2002). However a word can modify your prior polarity in relation to the context within which it is invoked. For example the word "earthquake" is used with negative meaning in the sentence :

"Selling the company caused an earthquake amount the employees".

Whereas it is used in an neutral meaning in the sentence:

"An earthquake is the result of a sudden release of energy in the Earth's crust that creates seismic waves".

For this reason, our system uses a coarse-grained WSD method for obtaining the contextual polarity of all words in tweets. The selected disambiguation method (Anaya-Sánchez et al., 2006) was developed for the traditional WSD task. In this WSD method, the senses are represented as topic signatures (Lin and Hovy, 2000) built from the repository of concepts of WordNet. The disambiguation process starts from a clustering distribution of all possible senses of the ambiguous words by applying the Extended Star clustering algorithm (Gil-García et al., 2003). Such a clustering tries to identify cohesive groups of word senses, which are assumed to represent different meanings for the set of words.

Resource	HP	HN	P	N	O
SWN	310	938	2242	2899	109035

Table 1: Senses highly positive, highly negative, positive, negative and objective distributions.

Then, clusters that match the best with the context are selected. If the selected clusters disambiguate all words, the process stops and the senses belonging to the selected clusters are interpreted as the disambiguating ones. Otherwise, the clustering is performed again (regarding the remaining senses) until a complete disambiguation is achieved. It does not distinguish between highly positive, positive, negative, highly negative or objective meaning of a word. In this paper, we propose a strategy to built a coarse-grained sense representation. Firstly, a topic signatures for all senses into WordNet is built and the topic signatures for coarse-grained senses is the sum of the topic signatures of the corresponding fine-grained senses that was grouped.

We explain coarse-grained sense representation using the following example:

Let us consider the adjective *"sad"*. This adjective has three word senses into WordNet 2.0

sad#a#1 – *experiencing or showing sorrow or unhappiness*
sad#a#2 – *of things that make you feel sad*
sad#a#3 – *bad; unfortunate*

Firstly the topic signature are built for each word sense:

$$vector1 = topicSignature(sad\#a\#1)$$
$$vector2 = topicSignature(sad\#a\#2)$$
$$vector3 = topicSignature(sad\#a\#3)$$

The senses are classified using equation (1)(in Section 2.2), sense 1 and 3 were considered as highly negative, whereas the sense 2 is objective. The topic signature associated to highly negative coarse-grained sense is computed as:

$$topicSignature(sad\#a\#HN) = sum(vector1 + vector3)$$

and objective coarse-grained sense is kept as vector2

$$topicSignature(sad\#a\#O) = vector2$$

2.4 Rule-based Sentiment Classifier

We use a rule-based classifier to classify tweets into positive, negative or neutral. A polarity value is assigned to each word, based on equation 2, after these were disambiguated. It is necessary to clarify that emotion words that replaced emoticons in the preprocessing phase, are not disambiguated. Instead, we give a prior polarity value equal to 4 if emotion word is *"happy"* and -4 in case that emotion word is *"sad"*. It is important to mention that the polarity of a word is forced into the opposite class if it is preceded by a valence shifter (obtained from the Negate category in GI (Stone et al., 1966)).

$$polarity(w) = \begin{cases} 4 & if\ w\ is\ disambiguated\ as\ HP \\ -4 & if\ w\ is\ disambiguated\ as\ HN \\ 2 & if\ w\ is\ disambiguated\ as\ P \\ -2 & if\ w\ is\ disambiguated\ as\ N \\ 0 & if\ w\ is\ disambiguated\ as\ O \end{cases} \tag{2}$$

The polarity of the tweet is determined from the scores of positive and negative words it contains. To sum up, for each tweet the overall positive ($PosS(t)$) value and overall negative value ($NegS(t)$) , are computed as:

$$PosS(t) = \sum_{w_i \in W_P} polarity(w_i) \tag{3}$$

W_P: Words disambiguated as highly positive or positive in tweet t

$$NegS(t) = \sum_{w_i \in W_N} polarity(w_i) \tag{4}$$

W_N: Words disambiguated as highly negative or negative in tweet t

If $PosS(t)$ is greater than $NegS(t)$ then the tweet is considered as positive. On the contrary, if $PosS(t)$ is less than $NegS(t)$ the tweet is negative. Finally, if $PosS(t)$ is equal to $NegS(t)$ the tweet is considered as neutral.

2.5 A Tweet Sentiment Classification Example

The general operation of the algorithm is illustrated in the following example:

Let us consider the following tweet:

@JoeyMarchant: I really love Jennifer Aniston :-) #loving, she is very coooolllll and sexy. I'm married to her... LOL, http://t.co/2RShsRNSDW

After applying the preprocessing phase, we obtain the following normalized text:

I really love Jennifer Aniston "happy" loving, she is very cooll and sexy. I'm married to her... lots of laughs.

When the text is lemmatized and stopwords are removed, we obtain the following bag of words (for each word we show: lemma and part-of-speech *n-noun, v-verb, a-adjective, r-adverb and u-unknown*):

really#r love#v jennifer#a aniston#n "happy"#a loving#a cooll#a sexy#a marry#v lot#n laugh#n.

After contextual word polarity detection, we obtain the following result (for each word we shown lemma, part-of-speech and sentiment sense, *HP-highly positive, HN-highly negative, P-positive, N-negative and O-objective*).

really#r#P love#v#P jennifer#a#O aniston#n#O "happy"#a loving#a#HP cooll#a#O sexy#a#P marry#v#O lot#n#O laugh#n#P

Once that all words were disambiguated we obtained their polarities using the equation 2 introduced in section 2.4. We show the polarities values assigned to each word, in Table 2.

Word	POS	Sentiment	Polarity
really	r	P	2
love	v	P	2
jennifer	a	O	0
aniston	n	O	0
"happy"	a	-	4
loving	a	HP	4
cooll	a	O	0
sexy	a	P	2
marry	a	O	0
lot	n	O	0
laugh	n	P	2

Table 2: Polarity assigned to each word

Note that the word "happy" has not been disambiguated, its polarity is assigned according to the emoticon associated in the original tweet. Afterward we compute overall positive and negative polarity value:

$$NegS(t) = 0$$

$$PosS(t) = 2+2+4+4+2+2 = 16$$

Therefore, the tweet t is classified as positive.

3 Results

This section presents the evaluation of our system in the context of SemEval 2013 Task No.2 Subtask B (Sentiment Analysis in Twitter). For evaluating the participant's systems two unlabeled datasets were provided, one composed of Twitter messages and another of SMS messages. For each dataset two runs can be submitted, the first (constrained), the system can only be used the provided training data and other resources such as lexicons. In the second (unconstrained), the system can use additional data for training. Our runs are considered as constrained because SSA-UO only use lexical resources for sentiment classification.

Runs	Dataset	F1	all runs Rank
twitter-1	Twitter	50.17	33(48)
sms-1	SMS	44.39	33 (42)

Table 3: SSA-UO results in polarity classification, all runs summited

Runs	Dataset	F1	constrained runs Rank
twitter-1	Twitter	50.17	25 (35)
sms-1	SMS	44.39	22 (28)

Table 4: SSA-UO results in polarity classification, constrained runs summited

In Table 3 we summarize the results obtained by SSA-UO system. As may be observed average F1 measure for Twitter dataset is the 50.17 and 44.39 for the SMS dataset. A total of 48 runs were submitted by all systems participant's in Twitter and 42 for SMS dataset. Our runs were ranked 33^{th} for both datasets.

In Table 4 we compare our results with those runs that can be considered as constrained. A total of 35 runs for Twitter and 28 for SMS were submitted , ours runs were ranked in 25^{th} and 22^{th} respectively. It's worth mentioning that, the results obtained can be considered satisfactory, considering the complexity of the task and that our system is unsupervised.

4 Conclusion

In this paper, we have described the SSA-UO system for Twitter Sentiment Analysis Task at SemEval-2013. This knowledge driven system relies on unsupervised coarse-grained WSD to obtain the contextual word polarity. We used a rule-based classifier for determining the polarity of a tweet. The experimental results show that our proposal is accurate for Twitter sentiment analysis considering that our system does not use any corpus for training.

Acknowledgments

This research work has been partially funded by the Spanish Government through the project TEXT-MESS 2.0 (TIN2009-13391-C04), "Análisis de Tendencias Mediante Técnicas de Opinión Semántica" (TIN2012-38536-C03-03) and "Técnicas de Deconstrucción en la Tecnologías del Lenguaje Humano" (TIN2012-31224); and by the Valencian Government through the project PROMETEO (PROMETEO/2009/199).

References

Balahur Alexandra, Steinberger Ralf, Goot Erik van der, Pouliquen Bruno, and Kabadjov Mijail. 2009. Opinion mining on newspaper quotations. In *Proceedings of the 2009 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology - Volume 03*, WI-IAT '09, pages 523–526, Washington, DC, USA. IEEE Computer Society.

Henry Anaya-Sánchez, Aurora Pons-Porrata, and Rafael Berlanga-Llavori. 2006. Word sense disambiguation based on word sense clustering. In *Proceedings of the 2nd international joint conference, and Proceedings of the 10th Ibero-American Conference on AI 18th Brazilian conference on Advances in Artificial Intelligence*, IBERAMIA-SBIA'06, pages 472–481, Berlin, Heidelberg. Springer-Verlag.

Stefano Baccianella, Andrea Esuli, and Fabrizio Sebastiani. 2010. Sentiwordnet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining. In Nicoletta Calzolari (Conference Chair), Khalid Choukri, Bente Maegaard, Joseph Mariani, Jan Odijk, Stelios Piperidis, Mike Rosner, and Daniel Tapias, editors, *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC '10)*, Valletta, Malta, may. European Language Resources Association (ELRA).

Andrea Esuli and Fabrizio Sebastiani. 2006. Sentiwordnet: A publicly available lexical resource for opinion mining. In *In Proceedings of the 5th Conference on Language Resources and Evaluation (LREC'06*, pages 417–422.

R. Gil-García, J. M. Badía-Contelles, and A. Pons-Porrata. 2003. Extended Star Clustering Algorithm. In *CIARP 2003, LNCS, vol. 2905*, pages 480–487.

Vasileios Hatzivassiloglou and Kathleen R. McKeown. 1997. Predicting the semantic orientation of adjectives. In *Proceedings of the eighth conference on European chapter of the Association for Computational Linguistics*, EACL '97, pages 174–181, Stroudsburg, PA, USA. Association for Computational Linguistics.

Jaap Kamps and Maarten Marx. 2002. Words with attitude. In *First International WordNet conference*.

Chin-Yew Lin and Eduard Hovy. 2000. The automated acquisition of topic signatures for text summarization. In *Proceedings of the 18th conference on Computational linguistics - Volume 1*, COLING '00, pages 495–501, Stroudsburg, PA, USA. Association for Computational Linguistics.

George A. Miller. 1995. Wordnet: A lexical database for english. *Communications of the ACM*, 38:39–41.

Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. 2002. Thumbs up? sentiment classification using machine learning techniques. In *Proceeding of Empirical Methods in Natural Language Processing*, pages 79–86.

Helmut Schmid. 1994. Probabilistic part-of-speech tagging using decision trees.

Philip J. Stone, Dexter C. Dunphy, Marshall S. Smith, and Daniel M. Ogilvie. 1966. *The General Inquirer: A Computer Approach to Content Analysis*. MIT Press, Cambridge, MA.

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

Michael Wiegand, Alexandra Balahur, Benjamin Roth, Dietrich Klakow, and Andrés Montoyo. 2010. A survey on the role of negation in sentiment analysis. In *Proceedings of the Workshop on Negation and Speculation in Natural Language Processing*, NeSp-NLP '10, pages 60–68, Stroudsburg, PA, USA. Association for Computational Linguistics.

Theresa Wilson, Janyce Wiebe, and Rebecca Hwa. 2006. Recognizing strong and weak opinion clauses. *Computational Intelligence*, 22:73–99.

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