

PIRPO: An Algorithm to Deal with Polarity in Portuguese Online Reviews from the Accommodation Sector

Marcirio Silveira Chaves¹, Larissa A. de Freitas²,
Marlo Souza², and Renata Vieira²

¹ Business and Information Technology Research Center (BITREC),
Universidade Atlântica, Oeiras, Portugal
marcirio.chaves@uatlantica.pt

² Faculdade de Informática,
Pontifícia Universidade Católica do Rio Grande do Sul, Porto Alegre, Brazil
{larissa.freitas,marlo.souza}@acad.pucrs.br, renata.vieira@pucrs.br

Abstract. This paper presents the algorithm Polarity Recognizer in Portuguese (PIRPO) to classify sentiment in online reviews. PIRPO was constructed to identify polarity in Portuguese user generated accommodation reviews. Each review is analysed according to concepts from a domain ontology. We decompose the review in sentences in order to assign a polarity to each concept of the ontology in the sentence. Preliminary results indicate an average F-score of 0.32 for polarity recognition.

Keywords: Portuguese Online Reviews, Portuguese Sentiment Analysis, Portuguese Opinion Mining, Accommodation Sector.

1 Introduction

Portuguese user generated reviews can be a valuable source of information for customer orientated business such as restaurants, hotels and airlines. However, manually processing lots of user generated content (UGC), such as reviews or tweets, is a costly and time-consuming task.

Sentiment in user generated content can be an indicator of wider customer opinion, consequently it is necessary to automatically recognise the sentiment expressed by the reviewers. According to Liu [7], sentiment classification is an area that has had a great deal of attention in the community. It treats sentiment analysis as a text classification problem, in which document-level sentiment classification aims to find the general sentiment of the author in an opinionated text. However, as document level sentiment classification fails to detect sentiment about individual aspects of the topic [10], in this work we separate out review attributes (concepts or features) within review and treat them differently instead of seeing a review as a whole, as suggested by Tang et. al. [10].

The objectives of this paper are two fold: (a) present the algorithm Polarity Recognizer in Portuguese (PIRPO) based on specific sets of handcrafted

lexico-syntactic patterns to sentiment classification in Portuguese online reviews and (b) evaluate the proposed algorithm using comments annotated by a human as reference for evaluation. We adopt the definition to sentiment given by Wilson et. al. [12]: “a positive or negative emotion, evaluation, or stance”.

2 Related Work

A popular Opinion Mining (OM) approach is a lexicon based strategy [9]. Lexicon-based approaches include the use of a list of nouns, verbs, adjectives and adverbs [4]. Ding et. al. [5] use all of these parts of speech in a holistic lexicon-based approach.

Chesley et. al. [4] use subjective and objective features, and polarity features to classify English opinionated blog texts. The first features include subjective (e.g. believing and suggesting) and objective (e.g. asking and explaining) verb classes, textual features (e.g. exclamation points and question marks) and Part-Of-Speech (POS) (i.e. first- and second-person pronouns and the number of adjectives and adverbs). Polarity features are composed by positive (e.g. approving and praising) and negative (e.g. abusing and doubting) verb classes, and positive and negative adjectives provided by Wikitionary. PIRPO uses: 1. subjective features, 2. objective features and 3. polarity features, but reviews are not annotated with a POS tagger.

Ding et. al. [5] determine the semantic orientations (i.e. polarities) of opinions expressed on product features in reviews. Their system processes reviews using a set of rules based on the polarity of the adjectives in each sentence. The algorithm uses the context of the previous or the next sentence (or clauses) to resolve opinion ambiguity. The idea is that people usually express the same opinion (positive or negative) across sentences unless there is an indication of opinion change using words such as “but” and “however”. Experiments indicate a F-score of 91 outperforming state-of-the-art methods. Instead of English, PIRPO is designed to analyse the polarity of service features in reviews written in Portuguese.

Kar and Mandal [6] use lists of adverbs and adjectives and apply fuzzy logic to find opinion strength in English web reviews. They also use a POS tagger and a list of stop words before detect features and extract patterns based on the POS tags. The intensity finder is the module where the features are weighted and the review is scored based on the adverbs and adjectives found in the review. As Kar and Mandal, we want to quantify the sentiment expressed in customer reviews into a score. While their system assigns a score to a review, PIRPO assigns a score to each concept of an ontology expressed in a review.

SentiStrength [11] is an emotion detection algorithm which uses a sentiment word strength list, composed by positive and negative terms classified for either positive or negative sentiment strength. This algorithm also uses a set of rules to correction spellings and punctuation and two lists: negative words and emotions. SentiStrength was applied to process MySpace comments, which tend to have more spelling errors and atypical punctuation and emoticons. PIRPO uses list of adjectives with the corresponding polarity (positive, negative or neutral)

associated to it. As we can find in the literature, a number of popular approaches for sentiment classification have been used to process English texts but only a small number have been adapted for Portuguese.

Recently, there have been a number of commercial sentiment tools available from software companies, for example: Attensity (<http://attensity.com>), Clarabridge (<http://clarabridge.com>), Sentimetrix (<http://www.sentimetrix.com>), and Synthesio (<http://synthesio.com>). Attensity and Synthesio explicitly claim to process Portuguese whilst Clarabridge and Sentimetrix are language independent, but there is no publicly available evaluation of these tools on Portuguese.

3 Our Approach to Recognise Polarity in Portuguese Online Reviews

Figure 1 presents the PIRPO information architecture. PIRPO receives as input a set of reviews which are pre-processed in order to extract their sentences and detect whose reviews that are split into positive and negative segments. Some information sources allow users to input positive and negative opinion in separately fields.




Fig. 1. PIRPO Information Architecture

PIRPO has a main module which recognise polarity of each sentence. This module receives as input sentences, a list of adjectives and concepts from an ontology. The output is a list of sentences with polarity recognised in each concept detected, which is stored in a knowledge base named Senti Knowledge Base.

The module Semantic Orientation Recognizer receives as input:

Reviews: The reviews are about Small and Medium Hotels in the Lisbon area, and it was provided by Chaves et. al. [2,3]. The information sources are Tripadvisor (www.tripadvisor.com) and Booking.com (www.booking.com). Full dataset is composed by 1500 reviews from January 2010 to April 2011 in Portuguese, English and Spanish, from which 180 in Portuguese.

Ontology Concepts: The concepts used to classify the reviews are provided by Hontology, which in its current version, has 110 concepts, 9 object properties and 30 data properties. To the best of our knowledge, Hontology is the first multilingual (i.e. Portuguese, English and French) ontology for the hotel domain. More detail about this ontology is given in Chaves et. al. [1].

List of adjectives: It is composed of sentiment-bearing words. This list of polar adjectives in Portuguese contains 30.322 entries provided by Souza et. al. [8]. This list is composed by the name of the adjective and a polarity which can assign one of three values: +1, -1 and 0. These values corresponding to the positive, negative and neutral senses of the adjective. PIRPO uses this list to calculate the semantic orientation of the concepts found in the sentences.

PIRPO assigns one out of three values to the polarity of a review segment: positive, negative, and mixed (when the segment contains one positive opinion and another negative about an object).

PIRPO output is a list of sentences with polarity that reflects the polarity of the words characterising the concepts of the ontology in the reviews. This approach eases the extension of PIRPO for other systems which use a different level of granularity.

```

Algorithm PIRPO(sentences, adjectives, concepts)
begin
  for each sentence  $s_i$  in sentences do
    sentence_polarity = 0
    for each concept  $c_j$  in concepts do
      for each word  $w_k$  in  $s_i$  do
        if  $c_j = w_k$  then
          let window as  $w_{k-n}$  to  $w_{k+n}$ , without  $w_k$ 
          for each adjective  $a_x$  in adjectives do
            if  $a_x$  is in window then
              sentence_polarity += polarity( $a_x$ )
            endif
          endfor
        endif
      endfor
    endif
    save ( $s_i$ ,  $c_j$ , sentence_polarity)
  endfor
endfor
end

```

PIRPO is implemented in Python. The algorithm above basically goes through each preprocessed sentence, each Hontology concept and each word in a pre-processed sentence. If the word in a pre-processed sentence is a concept of Hontology, we create a window of n words and we search these words in the list of adjectives. At least one word of the window should be in the list of adjectives for the sentence polarity to be calculated. The function *polarity* calculates the polarity based on the values in the list of adjectives. The function *save* saves

Table 1. Polarity assigned by a human evaluator

Concepts	Positive	Negative	Mixed
Hotel	3	2	22
Location	3	0	45
Room	3	5	28
Staff	2	2	35
Score	11	9	130

Table 2. Polarity classification by concept

Concepts	Precision			Recall			F-score		
	P	N	M	P	N	M	P	N	M
Hotel	0.20	0.00	1.00	0.33	0.00	0.22	0.25	0.00	0.36
Location	0.00	0.00	0.89	0.00	0.00	0.24	0.00	0.00	0.37
Room	0.11	0.00	1.00	1.00	0.00	0.19	0.20	0.00	0.32
Staff	0.08	0.00	1.00	1.00	0.00	0.12	0.15	0.00	0.21
Average	0.10	0.00	0.97	0.58	0.00	0.19	0.15	0.00	0.32

the results in comma separated values (CSV), which is to be stored in the Senti Knowledge Base.

4 Experiment, Evaluation and Discussion on the Results

The dataset selected to evaluate PIRPO is composed of reviews on Small and Medium Hotels in the Lisbon area, and it was provided by Chaves et al. [2]. We select only the 180 comments in Portuguese, which constitutes the subset used in our experiments. Each review sentence was evaluated as its polarity by a human. Polarity was classified as positive, negative or mixed. Table 1 presents the human evaluation.

According to Table 1, 86% of the human evaluations were classified as mixed. The concepts hotel, location, room and staff were the most mentioned in the online reviews. This fact indicates that these concepts were evaluated by hotel guests with positive and negative aspects. An in-depth analysis of the concepts classified as mixed is needed to find out more fine-grained information.

We evaluated PIRPO against the human evaluation. Table 2 presents the results for each concept. According to these results, PIRPO reached a better recall for concepts with positive polarity, while mixed polarity had a higher precision. We plan to verify the inter-annotator agreement using the kappa coefficient. Moreover, we also plan to evaluate PIRPO with other annotated corpus of online reviews.

5 Final Remarks

This paper introduced PIRPO, an algorithm to deal with polarity in Portuguese online reviews from accommodation sector and also described its preliminary

results. PIRPO was designed to be domain-independent and can be applied to other domains such as tourism or books. To do this we need other domain ontologies. Domain specific lexicon might also improve the classification task.

As future work, we will work on the improvement of the results reached by PIRPO so far. Moreover, we are going to extend PIRPO to recognize the strength of polarity underlying in each review. The final goal is to be able to integrate PIRPO in higher-level NLP systems, such as Information Extractors and Question Answer.

References

1. Chaves, M.S., Trojahn, C.: Towards a multilingual ontology for ontology-driven content mining in social web sites. In: 1st International Workshop on Cross-Cultural and Cross-Lingual Aspects of the Semantic Web (ISWC 2010), Shanghai, China (2010)
2. Chaves, M.S., Gomes, R., Pedron, C.: Analysing reviews in the Web 2.0: Small and Medium Hotels in Portugal. *Tourism Management* (2011)
3. Chaves, M. S., Gomes, R. e Pedron, C.: Decision-making based on Online Reviews: The Small and Medium Hotel Management. In: 20th European Conference on Information Systems, Barcelona, Spain, June 10-13 (to appear, 2012)
4. Chesley, P., Vincent, B., Xu, L., Srihari, R.: Using verbs and adjectives to automatically classify blog sentiment. In: AAI Symposium on Computational Approaches to Analysing Weblogs (AAAI-CAAW), pp. 27–29 (2006)
5. Ding, X., Liu, B., Yu, P.S.: A holistic lexicon-based approach to opinion mining. In: Proceedings of First ACM International Conference on Web Search and Data Mining, WSDM 2008 (2008)
6. Kar, A., Mandal, D.P.: Finding opinion strength using fuzzy logic on web reviews. *International Journal of Engineering and Industries* 2(1) (2011)
7. Liu, B.: Sentiment analysis and subjectivity, 2nd edn. *Handbook of Natural Language Processing* (2010)
8. Souza, M., Vieira, R., Busetti, D., Chishman, R., Alves, I.M.: Construction of a Portuguese opinion lexicon from multiple resources. In: The 8th Brazilian Symposium in Information and Human Language Technology (STIL 2011), Cuiabá, Brazil (2011)
9. Taboada, M., Brooke, J., Tofiloski, M., Voll, K.D., Stede, M.: Lexicon-based methods for sentiment analysis. *Computational Linguistics* 37(2), 267–307 (2011)
10. Tang, H., Tan, S., Cheng, X.: A survey on sentiment detection of reviews. *Expert Systems with Applications* 36(7), 10760–10773 (2009)
11. Thelwall, M., Buckley, K., Paltoglou, G., Cai, D., Kappas, A.: Sentiment in short strength detection informal text. *JASIST* 61(12), 2544–2558 (2010)
12. Wilson, T., Wiebe, J., Hoffmann, P.: Recognizing contextual polarity: An exploration of features for phrase-level sentiment analysis. *Computational Linguistics* 35, 399–433 (2009)