

Determining the Level of Clients' Dissatisfaction from their Commentaries

Ana Catarina Forte¹ and Pavel B. Brazdil^{1,2}

¹LIAAD – INESC Tec, Porto, Portugal

²FEP, Univ. of Porto, Portugal

Abstract. We present a study in the area of sentiment analysis of clients' commentaries transcribed by assistants of a help-desk service of one Portuguese telecommunications company. We have adopted a lexicon-based approach to determine the polarity of the sentiment of each commentary, based on the so called *opinion words*. This task was by no means easy, as not many tools are available for the Portuguese language. The initial results with the off-the-shelf solutions were rather poor. This has motivated us to carry out a number of enhancements, including, for instance, enriching the given lexicon with *domain specific* terms, formulating specific rules for negation and amplifiers. Automatic pruning of some of the lexicon terms has led to a significant improvement in performance. As our final system achieved a very good performance, our work should be of interest to others working on domain specific solutions for languages where ready-made solutions are not available.

Keywords: Sentiment analysis in Portuguese, Lexicon-based approaches, Lexicon enrichment, Negation of sentiment words, Amplifiers and attenuators.

1 Introduction

In recent years characterized by emergence of social networks, blogs and forums, individuals and organizations used these means to express their opinions. As a result, there has been a continuous increase in the daily volume of information generated by users of social platforms, which created new challenges. First, it is not humanly possible to read all this information. Another problem is that organizations may be unaware of what the customers' opinions are. The customers may publish their opinions, sometimes quite negative ones, which are disseminated quickly reaching a large number of people and the company may be unaware of this. Thus organizations strive to continuously update their information, so as to understand the customers' behavior better, enabling them to predict their next steps.

According to some sources [3], 80% of customer data captured within enterprises is in unstructured format, in the form of call center transcriptions or e-mails from

customers and employees. It is relevant to study how customers manifest themselves over time, enabling to detect changes in their preferences and identify patterns of their behavior or detect that customers may effectively stop paying for certain services or products.

This situation has motivated us to undertake this study. The data comes from one customer service department of a major Portuguese telecommunications company. This service involves questions concerning customer satisfaction, promotional campaigns, password change and description of problems that need to be resolved. It provides means for the customers to express their opinions and ask questions. This service is coordinated by technical assistants who answer phone calls from customers, while transcribing the contents of each call in a form of *commentaries*.

Our aim was to analyze this data and determine in an automatic way the polarity of each commentary. We have adopted a lexicon-based approach and employed two lexicons (*Afinn*, *Sentilex*) initially, but only *Sentilex* could be used directly with Portuguese texts. The terms of *Afinn* had to be translated from English into Portuguese. Unfortunately, the results were rather disappointing. This has motivated us to carry out various enhancements whose aim was to improve the results. One involved the inclusion of *domain specific* terms in the lexicon which lead to a marked improvement in performance. We have also formulated specific rules determining how negation should be treated in Portuguese, and similarly, investigated also *intensifiers* and *attenuators*. Each of these led again to improved results. Perhaps the most surprising finding was that the method of automatic pruning of some of the lexicon terms proved beneficial for performance. The performance of the final system was indeed very good in terms of predictive accuracy of (81.9%) and also, in terms of cost-sensitive analysis described later.

The structure of the rest of this paper is as follows: Section 2 presents a review of the literature in the area of sentiment analysis and special attention is paid to works focusing on Portuguese texts. Section 3 starts with an overview of a typical sentiment analysis system and then continues with the description of different enhancements of the system that led to improved results. Section 4 discusses the results obtained and Section 5 presents conclusions and suggestions for future work.

2 Related Work

Sentiment analysis (SA) is a field of study that examines opinions, feelings, evaluations, attitudes and emotions in relation to entities such as products, services, organizations, people, events and topics [10]. Most work in this area is carried out for the English language [12, 13]. For the Portuguese language the available resources are rather scarce. As it has rather specific characteristics (e.g. multiple verbal forms, specific rules for negation), developing a SA system can be considered a challenge.

Several methods have been proposed for the automatic determination of a sentiment value in text documents [10, 13, 17]. The most common approach involves considering this a text classification task, where the aim is to determine whether a document expresses a *negative*, *neutral* or *positive* opinion. The two main approaches are a lexicon-based approach and a machine learning approach [22].

The machine learning approach uses a set of features to learn a classifier. The features represent *terms* or in general *unigrams*, *bigrams* or *multi-word units*, which may in addition be characterized by part-of-speech tags. Although this approach may obtain good results, the transfer of the model - trained classifier - to another domain may not lead to good results [2].

The lexicon-based system can be developed manually [21], [23] or automatically [6], [24], [25]. This approach uses a *sentiment lexicon*, which includes words or expressions together with their polarities. It is important that the lexicons have a good coverage of words that appear in the texts that need to be processed.

As for the number of polarity levels, some researchers have considered just three (e.g. *negative*, *neutral* and *positive* or $-1,0,1$ if we use numbers), while other used more. In the work reported here we use 7 levels ($-3,-2,-1,0,1,2,3$), where the value -3 (3) represents the most negative (positive) sentiment value. We have adopted 7 levels because our client – the telecommunications company – has requested this.

Some authors claim that lexicon-based methods are more robust than the machine learning approaches (ML) [22]. However, our initial experiments with the lexicon-based approach did not result in good performance in our application domain. This indicates that this approach may also suffer from fragilities attributed by Taboada et al. [22] to ML approaches.

One study in Portuguese [18] determines the sentiment of texts from Twitter messages relative to five political leaders involved in the 2011 parliamentary elections. The task was first to identify the political leaders and then determine the polarity of each Twitter message. The sentiment analysis was based on lexical polarity, lexical-syntactic patterns and inference rules.

Some people have also studied the so called *polar expressions*. Polanyi et al. [15] have proposed the following rule to deal with such expressions: “*If a given polar expression includes negation, the sentiment value of the associated expression should be reversed*”. This rule is also applicable to Portuguese texts, although as we will see further on, in some situations it is useful to use a somewhat different rule.

Another line of work involves considering *amplifiers* and *attenuators* of the sentiment value. The most common way to identify these terms is by a recourse to a list of *adverbs* and *adjectives*. Polanyi et al. [15] modified the sentiment value of the subsequent word by adding or subtracting 1. For example, suppose the amplifier *very* has been used within the expression *very beautiful*. As *beautiful* has a sentiment value of $+2$, the resulting value of the expression is $2+1=3$.

Brooke [1] has proposed to treat *amplifiers* in a somewhat different way. Instead of adding (or subtracting) a fixed value, the author proposed to increase (or decrease) the sentiment value by a certain percentage. However, the author did not provide details on how the percentage is determined.

3 Adapting the Approach of Sentiment Analysis for Portuguese

As some readers may not be familiar with sentiment analysis (SA), the next subsection provides an overview of a typical lexicon-based approach. This serves also as a basis for describing the enhancements carried out for Portuguese (subsection 3.2).

3.1 Overview of a Typical Lexicon-based Approach to Sentiment Analysis

The typical lexicon-based sentiment analysis method involves basically two major steps, pre-processing and calculation of the sentiment value of texts using sentiment lexicons. Both are described in more detail below.

Pre-processing

Many pre-processing steps used in SA tasks are commonly used also in classification of texts or clustering. Therefore our description here is rather brief, as readers may consult other sources [4, 7, 20].

One important step involves transformation of capital letters to lower case enabling to identify words that are similar, but not written in exactly the same way. The following steps normally include removal of numbers, punctuation symbols, short words (e.g. consisting of 1-2 characters), blanks and *stopwords*. All these items are normally not relevant and their elimination does not usually effect the outcome.

Many systems apply also *stemming* which consists of eliminating the suffixes, while retaining only the root of the word. Stemming algorithms depend strongly on word formation rules. So, direct application of an English *stemmer* may result in an undesirable outcome when applied to other languages including Portuguese [8]. This is why some researchers have adopted *lemmatization* instead, which selects a representative form (e.g. infinitive) from a given set of words with the same root.

Usage of lexicons to calculate the sentiment value

After the preprocessing has been carried out, the chosen sentiment lexicon is used to derive the sentiment value of a given text. This is done by processing all words. If a particular word appears in the sentiment lexicon, its sentiment value is recorded. The sentiment value of the given text is determined by summing up all the sentiment values encountered. This process is illustrated by the following example:

cliente quer (2) desativar (-2) serviços pois vai sair sem período fidelização

(the customer wants to stop the services, because he will leave without the loyalty period)

The only words that appear in a sentiment lexicon are *quer* and *desativar* with sentiment values of +2 and -2. So the resulting sentiment value of this phrase is 0.

3.2 Enhancements of the Basic Approach for Sentiment Analysis in Portuguese

The results of the basic system just described, when applied to our document set in Portuguese were rather disappointing, because the predictive accuracy was rather low. This is the reason why we have set out to carry out different enhancements with the objective to improve of performance, acceptable to the telecommunications company.

3.2.1 Our Dataset and the Lexicons Used

The document collection use here consists of commentaries of clients that were gathered in a period of about 3 months in 2014, totaling 1724 cases. Each document includes one commentary, consisting of about 17 words on average.

Lexicon *Sentilex*: This sentiment lexicon was elaborated for the Portuguese language and contains more than 82k entries. It was developed specifically for data mining applications and classification of sentiments in Portuguese [19]. The entries can be *inflected forms of words*, which can be included within expressions. Expressions can contain a one or more words (*compound expressions*) and typical idioms. The sentiment values are on an integer scale ranging from -1 to 1.

Lexicon *Afinn* and *AffinP*: Lexicon *Afinn* includes 2477 words and phrases in English, whose sentiment value range from -5 to 5, where -5 (5) is associated with very negative (positive) words [11]. This lexicon was translated into Portuguese manually which was quite time consuming task, lasting about 3 weeks. The Portuguese version, *AffinP*, includes 1997 words.

Domain specific lexicon *AffinPD*: We have noted that the lexicons above did not include many of the terms that were used in clients' commentaries. So, for instance, the words *comprar* and *fidelização* (*buy* and *loyalty agreement*) which do not exist in *Sentilex*, nor in *AffinP*, are indicative of a positive action on the part of client. The word *reclama* (*complains*) exists only in *AffinP*, but not in *Sentilex*. It transmits clearly a negative sentiment.

This analysis has motivated us to explore the hypothesis that elaborating a domain specific lexicon may lead to a marked improvement of performance. We have decided to use *AffinP* as a starting point, as this lexicon attained better performance than *Sentilex*. Some of sentiment values in *AffinP* were altered, taking into account the specificities of the domain of telecommunications. The domain specific lexicon is referred to as *AffinPD* and includes 6915 words. It uses the sentiment scale of *Afinn* ranging from -5 to 5. As this lexicon was elaborated manually, this task was rather time consuming and required about 2 months to complete.

The percentage of words covered by *AffinPD* of our text documents was around 88.3%, whereas in the other lexicons this rate was quite low (reaching about 20%). Our experimental results presented later show that the addition of domain specific terms leads to a marked improvement of performance.

Optimizing the Lexicon by Pruning: It is conceivable that not all terms in the given lexicon are useful and that maintaining them there may in fact be detrimental to performance. We have decided to investigate this issue. This study was conducted only for *AffinPD* lexicon. The method used for constructing the optimized lexicon is similar to the feature selection method called *Backward Elimination* [4]. The method processes all terms in turn. In each step the system determines the effect of eliminating a particular term on the overall performance of the system. If the performance decreased, the term is maintained, but otherwise it is be dropped. In this phase we have used a validation set (different from the test set) which consisted of about 600 commentaries accompanied by a sentiment value.

3.2.2 Different Enhancements Carried Out

Enhanced Pre-Processing

The pre-processing steps described in section 3.1 were re-used in the system described in this section, which includes several enhancements described next.

Restricted Stemming: It may happen that some words that have different sentiment value before stemming become indistinguishable after stemming. Some examples are shown in Table 1.

Table 1. Undesirable effects of stemming for some Portuguese words

Lexicon	Word	Sentiment value	Root
AffinP	<i>afetado</i> (affected)	-2	afet
AffinP	<i>afeto</i> (feeling)	3	afet
Sentilex	<i>fala</i> (speaks)	0	fal
Sentilex	<i>falido</i> (bankrupt)	-1	fal

For this reason, in this work stemming is *applied selectively* only to words that do not appear in the sentiment lexicon. Usage of lemmatization would have been another possible alternative.

Rewrite rules: The texts written by operators often include several different variants of the same word (e.g. *desativar*, *desactivar*). This is undesirable, as some of these variants may not appear in the lexicon. Therefore, we have decided to formulate a set of rules (635 in total) that substitute certain abbreviations and misspelled variants by their standardized form. Below we give two examples of such rules:

clt → cliente
desactivar → desativar

The second rule transforms the word *desactivar* from the old orthographic convention to the new one.

Removal of short texts: Our enhanced preprocessing comprises also the removal of commentaries which include only 1 word, as these transmit hardly any information.

Domain-specific stopwords: This stage focuses on building a list of domain-specific *stopwords*, or predefined phrases recorded by operators that do not transmit any unusual sentiment value. This list includes 345 items. Consider, for instance:

Boa tarde, bem-vindo ao apoio ao cliente se pretende ajuda sobre um tema (...)
Good afternoon, welcome to customer support if you want help on a topic (...)

The words *boa* (good), *apoio* (support), *ajuda* (help) would be attributed a positive sentiment value and hence influence the final scores. However, this does not make sense, as these words do not convey any sentiment value on the part of the client.

3.2.3 Dealing with Polar Expressions

Our system deals also with expressions that include *negation* which are sometimes referred to as *polar expressions*. These may include, for instance, a negation word, such as *não* (*no* or *not* in English) followed by some other word or expression, such as *divertido* (*funny* in English). If we did not introduce any special treatment, we could obtain a counterintuitive result. Suppose the sentiment value of *não* is -1 (as in *AffinP*) and the sentiment value of *divertido* is 3, the resulting value would be positive.

The rule proposed by Polanyi et al. [15]: “*If a given polar expression includes negation, the sentiment value of the associated expression should be reversed*”, suggests a solution, which is also applicable to Portuguese texts. So, in our example with *não* (-1) *divertido* (3), the resulting value would be $-1 \times 3 = -3$. In the following this rule is referred to as *inversion rule* and is used in all cases where the word appearing after negation has a positive sentiment.

However, there are cases when the word appearing after negation has a negative sentiment. The application of the above rule would lead to an incorrect value. Consider, for instance, the phrase “*cliente não* (-1) *quer ser mais incomodado* (-2)” (customer *doesn't* want to be more *disturbed*). It does not seem right to attribute it a positive sentiment value resulting from $-1 \times (-2) = 2$. In our system this phrase is attributed the value of -2 (the sentiment value of *incomodado*). We will call this rule *corrected inversion rule* in the following. The results presented in Section 4 indicate that the *corrected inversion rule* is preferable to the original rule.

Our system permits to process various other words besides the negation word *não*. The list assembled contains 33 different words that modify the meaning of the subsequent text. The list includes, for instance *não*, *nunca*, *jamais*, *nem*, *ninguém*, *nenhum*, *negativamente*, *tampouco*, among others.

3.2.4 Amplifiers and Attenuators

Polanyi et al. [14], define *contextual valence shifters* as contextual phenomena that modify the polarity of another term. According to the authors these can be divided into three groups: *inverters* invert the polarity of a polarized item, *intensifiers* (or *amplifiers*) intensify it and *attenuators* diminish it [9, 16]. Sentiment modifiers are mostly *adverbs* (*attractively*, *annoyingly*), *adjectives* (*better*, *bad*), *sentence connectors* (*although*, *however*, *but*), *modal verbs* (*might*, *could*, *possibly*) or even *nouns* (*freedom*, *bankruptcy*).

We have followed this line of work, but adapted it to Portuguese, while focusing mainly on sentiment modifiers that are *adverbs* and *adjectives*. Words such as *demasiado* (*too much*), *bastante* (*in sufficient number*), *pouco* (*a little*) and *menos* (*less*) do not exist in the sentiment lexicons analyzed. To resolve this issue, we have assembled a list of 166 amplifiers and 35 attenuators. The process for calculating the sentiment value of sentences is as follows:

If an *amplifier* is encountered:

If the value of the subsequent text is positive (negative), add (subtract) 1.

If an *attenuator* is encountered:

If the value of the subsequent text is positive (negative), subtract (add) 1.

3.2.5 Normalizing the Sentiment Values

As the ranges of the sentiment values in the lexicons used are different, the resulting sentiment value of the given text (commentary) attributed by our system could depend on which lexicon has been used. This is obviously undesirable. We have therefore incorporated a normalization step, whose aim to transform the sentiment values generated into a normalized range spanning from -3 till 3. This range was suggested by the company who require that we not only identify the negative (positive) cases, but distinguish how grave the situation is.

Normalization is relatively simple to carry out. Basically, it is necessary to divide all values by normalization constant $N_{AFFINPD}$ or $N_{SENTILEX}$, depending on which lexicon is being used. After some analysis we have determined that the setting $N_{AFFINPD}=4$ and $N_{SENTILEX}=1$ is the most appropriate.

The sentiment values are illustrated by the following example that was already used in Section 3.1. Here we show the effects of using the enriched lexicon, negation rules and normalization.

cliente quer desativar (-5) serviços pois vai sair (-3) sem (-1) período fidelização (3)

Let us examine the effects of negation in this example. If we did not use it, the final sentiment value (before normalization) would be $-5-3-1+3=-4$. The inversion rule for negation transforms the last 3 words into *sem período fidelização (-3)* and hence the score is $-5-3-3=-11$. This value needs to be normalized, which leads to $-11/N_{AFFINPD} = -11/4 = -2.75$. Finally, the sentiment value is mapped to the nearest integer value -3.

3.3 Evaluation of the sentiment classifier

The data for the evaluation consisted of 200 texts accompanied by their sentiment value ranging between -3 and 3. The evaluation was carried out by comparing the sentiment values generated with the ones that were attributed by a person (ground truth). Two different performance measures were used - *success rate* and *mean cost per case*. The use of cost-sensitive analysis was motivated by the fact that the sentiment values -3, -2, .. 3 are on ordinal scale. It is therefore useful to consider how far a particular prediction is from the correct value. Assuming that -3 and 3 are one unit apart, the distance between -3 and -2, for instance, is a fraction of that distance, i.e. $1/6=0.17$ (after rounding). So this error has the cost of 0.17. The mean cost per case is calculated by calculating the mean of the costs of individual documents.

4 Results

The results presented here are relative to our test dataset consisting of about 200 cases accompanied by a sentiment value. The evaluation involved the set of 13 different variants of our sentiment analysis system described next. We have conducted experiments to determine to what extent the different facets incorporated in these variants would improve the performance. Fig.1 shows the results.

1. Lexicon *AffinP*: This variant represents the very basic SA system with *AffinP*. It achieved a modest success rate of 33.33% and mean cost of 16.85%.
 2. Lexicon *AffinPD*: Similar to variant 1, but the system uses an enriched lexicon containing domain-specific words. The success rate improves by 5.9%.
 3. Pre-processing: This variant extends the previous variant and uses the basic pre-processing method described in Section 3.1.
 4. Rewrite rules, as described in Section 3.2.2. This led to quite significant reduction of the mean cost by 1.88%.
 5. Removal of texts with a single word (see Section 3.2.2).
 6. Elimination of domain-specific stopwords (see Section 3.2.2).
 7. Elimination of stopwords: We have reused the list of 203 Portuguese stopwords available in R with the command `stopwords('portuguese')`.
 8. Restricted stemming (see Section 3.2.2).
 9. Negation 1: This variant incorporates the *inversion rule* (see Section 3.2.3)
 10. Negation 2: This variant extends the variant 8 by incorporating the *corrected inversion rule* (see Section 3.2.3). This led to a significant cost reduction of 3.87%.
 11. Amplifiers (see Section 3.2.3): This variant extends the variant 10. This led to a cost reduction of 1.25%.
 12. Attenuators (see Section 3.2.3).
 13. Optimized Lexicon *AffinPDO*: The lexicon is pruned following the method described in Section 3.2.1. The final version had a very good mean cost of 4.54%.
- Each of the following versions extends its predecessor unless stated otherwise.




Fig. 1. Results of different variants of our system

We have carried out also experiments with *Sentilex* lexicon which is available for Portuguese texts. Unfortunately, this system had poor results. The basic solution achieved the success rate of 31.34% and the mean cost of 18.1% (i.e. worse than our version 1 with *AffinP*). This lexicon had rather poor coverage of the words (only

about 5%) that appear in our commentaries. Also, the range of sentiment values range from -1 to 1, which is probably insufficient to capture the nuances of sentiment.

5 Conclusions

Our aim was to develop a sentiment analysis system which could be applied to Portuguese texts. This appeared a challenge, as not many tools were available. Applying off-the-shelf solutions in conjunction with lexicon Sentilex, available for Portuguese, gave poor results. We conjecture that this is due mainly to the fact that its terms do not cover the domain of telecommunications well, and also, the range of just three sentiment values is not able to capture the varied degrees of sentiments.

We have therefore analyzed a set of techniques with the objective of determining whether these could lead to improved performance. Translating terms of the chosen lexicon (*Affin*) into Portuguese and enriching it with domain-specific terms was rather time consuming, but the effort paid off. It led to a marked improvement of performance. There are many specific terms in the telecommunications that were not present in the original lexicon. Defining rules enabling to translate words into its standardized form was also useful.

We have also incorporated rules that deal with negation and adapted the inversion rule described in literature to Portuguese. We did not limit ourselves just to the negation word *não*, but defined an extended set of 33 *negation words*. We have investigated situations when the word following the *negation word* has a negative sentiment. The experiments on pre-classified data indicated that the *corrected negation rule* is preferable to the usual *inversion method* (as in [15]). Finally, we have incorporated rules that modify the sentiment value of expressions that include *amplifiers* and *attenuators*, which required that a list of words be built for the Portuguese language. The final system reached an excellent performance both in terms of accuracy and mean cost.

Future work may involve the incorporation of an automatic spell-checker for Portuguese. Elaborating rules manually to correct abbreviated or misspelled words is time-consuming. We could also investigate the possibility of generating the enriched lexicon through automated methods, rather than through a manual approach. As this task is rather time-consuming, a great deal of time could be saved. Also, it would make our system adaptable to other domains.

Acknowledgements

We wish to thank Nuno Paiva for defining this very interesting problem in the first place and to João Cordeiro for his useful comments on this work.

We acknowledge the support of *Project NORTE-01-0145-FEDER-000020*, financed by the *North Portugal Regional Operational Programme (NORTE 2020)*, under the *PORTUGAL 2020 Partnership Agreement and support from ERDF – European Regional Development Fund* through the *Operational Programme for Competitiveness and Internationalisation - COMPETE 2020 Programme* within project «*POCI-01-0145-FEDER-006961*», and by *National Funds* through the *FCT – Fundação para a Ciência e a Tecnologia (Portuguese Foundation for Science and Technology)* as part of project *UID/EEA/50014/2013*.

References

1. Brooke, Julian. 2009. *A Semantic Approach to Automatic Text Sentiment Analysis*. M.A. thesis, Simon Fraser University, Burnaby, B.C., Canada.
2. Gamon, M. and Aue, A. (2005). "Automatic identification of sentiment vocabulary: exploiting low association with known sentiment terms". *Proc. of the ACL WS on Feature Engineering for ML in Natural Language Processing*. Ann Arbor, Michigan, ACL, 57-64.
3. Group, F.C. (2010). Text Mining - Going Way beyond Just Listening to the Voice of the Customer. Forte Consultancy.
4. Guyon, I., Gunn, S., Nikravesh, M. and Zadeh, L. A. (2006). "Feature Extraction: Foundations and Applications", In *Studies in Fuzziness and Soft Computing*, Springer, NY.
5. Haddi, E., Liu, X. and Shi, Y. (2013). "The Role of Text Pre-processing in Sentiment Analysis", *Procedia Computer Science*, Vol. 17, pp. 26–32.
6. Hatzivassiloglou, V. and McKeown, K. R. (1997). "Predicting the Semantic Orientation of Adjectives", In *Proceedings of the 35th Annual Meeting of the ACL and 8th Conference of the European Chapter of the ACL*, pp.174-181, Madrid.
7. Hemalatha, I., Saradhi-Varma, GP. and Govardhan, A. (2012) "Preprocessing the Informal Text for efficient Sentiment Analysis", In *International Journal of Emerging Trends and Technology in Computer Science* , Vol. 1, Issue 2, pp. 58-61.
8. Honored, A., Leon, R., O'Donnel, R. and Sinclair, D. (2000). "A word stemming algorithm for the Spanish language". *Proc. of the Seventh Int. Symposium on String Processing Information Retrieval (SPIRE'00)*. IEEE Computer Society Washington, DC, USA, 139.
9. Kennedy A. and Inkpen D. (2006). "Sentiment classification of movie reviews using contextual valence shifters". In *Computational Intelligence*, 22(2):110–125.
10. Liu, Bing (2012). *Sentiment Analysis and Opinion Mining: Synthesis Lectures on Human Language Technologies*. Morgan & Claypool Publishers.
11. Nielsen, F. (2011). "A new ANEW: Evaluation of a word list for sentiment analysis in microblogs", *Proceedings of the ESWC2011 Workshop on Making Sense of Microposts: Big things come in small packages 718, CEUR Proceedings*, pp. 93-98.
12. Pang, B. and L. Lee (2008). "Opinion Mining and Sentiment Analysis." *Found. Trends Inf. Retr.* 2(1-2): 1-135.
13. Pang, B; Lee, L, and Vaithyanathan, S. (2002). "Thumbs up?: sentiment classification using machine learning techniques". *Proc. of the ACL-02 Conf. on Empirical methods in natural language processing – Vol. 10 (EMNLP '02)*. ACL, pp. 79-86.
14. Polanyi L. and Zaenen A. (2004). "Contextual valence shifters". In *Proceedings of AAAI Spring Symposium on Exploring Attitude and Affect in Text*, pp 106–111.
15. Polanyi, L. and Zaenen, A. (2006). "Contextual Valence Shifters". In *Computing Attitude and Affect in Text*, James G. Shanahan, Yan Qu, Janyce Wiebe (eds.), Springer Verlag, Dordrecht, The Netherlands, Vol. 20, pp. 1-10
16. Quirk, R., Greenbaum, S., Leech, G. and Svartvik, J. (1985). *A comprehensive grammar of the English language*, London: Longman.
17. Saraswat, M. and Patel, R. (2014). "A Survey on Sentiment Analysis". *Int. Journal of Research in Engineering Technology and Management*, Vol. 2, Issue: 6.
18. Silva, M.J. and Team, R. (2011). "Notas sobre a realização e qualidade do twitómetro". Technical report, University of Lisbon, Faculty of Sciences, LASIGE, Lisbon, Portugal.
19. Silva, M.J., Carvalho, P. and Sarmento, L. (2012). "Building a Sentiment Lexicon for Social Judgement Mining". In *International Conference on Computational Processing of the Portuguese Language (PROPOR)*, LNSC, Springer, pp. 218-228.
20. Soman, K.P., Diwakar, S. and Ajay, V. (2006). *Data Mining: Theory and Practice*, PHI Learning.
21. Stone, P.J., D.C. Dunphy, M.S. Smith, and D.M. Ogilvie (1966). "The General Inquirer: A Computer Approach to Content Analysis". *MIT Press*, Cambridge, MA.

22. Taboada, M., Brooke, J., Tofiloski, M., Voll, K. and Stede, M. (2011). "Lexicon-based methods for sentiment analysis". *Comput. Linguist*, 37, 2 267-307.
23. Tong, R.M. 2001. "An operational system for detecting and tracking opinions in on-line discussions". In *Working Notes of the ACM SIGIR 2001 Workshop on Operational Text Classification*, pp 1–6, NY.
24. Turney, P. (2002). "Thumbs up or thumbs down? Semantic orientation applied to unsupervised classification of reviews". In *Proceedings of 40th Meeting of the Association for Computational Linguistics*, pp 417–424, Philadelphia, PA.
25. Turney, P. and Littman, M. (2003). "Measuring praise and criticism: Inference of semantic orientation from association". *ACM Transactions on Inform. Systems*, 21(4):315–346.