Sentiment Analysis of Text Documents

Dalibor Bužić
College for Information Technologies
Klaićeva 7, Zagreb, Croatia
dalibor.buzic@vsite.hr

***Abstract**. By popularizing Web 2.0 technologies, the internet has become a rich source of user-generated content. Sentiment analysis collects and processes people's opinions and attitudes toward products, services, politics, social events, marketing campaigns and company strategies. In the sentiment analysis, unstructured text is processed, and it brings numerous problems in computer processing. In order to overcome problems, different methods and techniques are used. This paper gives a literature overview in the field of sentiment analysis focusing on current research approaches - machine learning, lexicon-based approach, dataless classification and deep learning. Emotion mining, a relatively new subfield within sentiment analysis, is also discussed in this paper.*

Keywords. sentiment analysis, opinion mining, machine learning, deep learning, dataless classification, emotion mining

1 Introduction

Over the last decade, the interest in collecting and processing text from social networks, forums, blogs, review sites and similar sources of user-generated content has been increasing. Text contains large amount of opinions of people with different profiles (age, region in which they live, education, viewpoints ...) about perception of products and services, political decisions, companies etc. Clearly, companies and political parties have had the same need in the pre-Internet age (which was solved by surveys, questionnaires and other ways), but by popularizing Web 2.0 technologies users have started to express their opinions publicly in digital form.

Due to the large amount of user content, a need for its efficient computer processing emerged very quickly. The first problem is that the text often contains spelling mistakes (mostly unwittingly, sometimes intentionally), abbreviations, new words, jargon, and other elements that make it difficult to process it. One field of computer science, natural language processing deals with this, as well as numerous other issues (word spelling, word stemming, entity recognition, machine translation, segmentation of texts by topic ...). One of the areas of application within the natural language processing is a sentiment analysis (in the past, the expression *opinion mining* was used more frequently).

The sentiment analysis is an area that analyzes the opinions, feelings, estimates and attitudes of people towards entities and their attributes in textual data (Liu, 2015, p.1). Although the sentiment analysis can be performed in image recordings (Xu et al., 2014) or in combination of image and text (Vadicamo et al., 2017), only textual content will be discussed in this paper.

Roughly speaking, sentiment analysis can be seen as a process that consists of the sequence of following steps: definition of purpose and setting goals, followed by retrieval of data, preprocessing, feature extraction, text classification according to sentiment, interpretation of the results and, at the end, their presentation.

Classification is usually done according to polarity or orientation. It is a two-step process (two binary classifications): text is first classified according to objectivity (the objective text does not contain sentiment and is eliminated from further analysis) after which it is classified into a positive or negative category. Alternative approach is to divide text into one of three classes (positive / neutral / negative) in one step. However, instead of classifying a subjective text only in a positive or negative class, sentiment analysis can also be carried out according to specific emotions such as surprise, fear, disgust, joy and trust.

The unit of the text over which the analysis is conducted can be:

- Document – analysis resolves whether a positive or negative opinion predominates about subject to which the document refers (e.g. film review),
- Sentence - the sentence can be viewed as a small document, and therefore there is no fundamental difference between the sentiment analysis at the document level and at the sentence level,
- Aspect - the analysis is based on some important properties or parts of the entity; it is common in reviews of restaurants (in the sentence *Food is excellent, but the service is terribly slow*, aspects are *food* and *services*) or products (in the sentence

Tires are noisy and expensive aspects are *volume* and *price*, although not explicitly stated).

Words are the main carriers of sentiment, but phrases can be as well. In addition to the basic classification of the word in positive, negative, or neutral, it is possible to assign numeric values to each of these classes, as it is done in the SentiWordNet dictionary where one word gets three decimal values whose total sum is 1 (Esuli & Sebastiani, 2006).

2 Applications of Sentiment Analysis

People had relied on the recommendations in the past, but were mostly limited to the opinions of friends and relatives, while today they are able to read a large number of opinions of unknown persons. Companies are aware of this, and sentiment analysis helps them understand how the reputation of their (as well as competing) product changes over time or, for example, the greatest weaknesses of the product. Between 73% and 87% of people claim that restaurant reviews, hotels and travel agencies have had a significant impact on the purchase; also the buyers were willing to pay from 20% to 99% more (the percentage depends on the type of product or service) for the product/service rated five-star than the one with the four stars (Pang & Lee, 2008).

The results of the sentiment analysis can also be used for social issues. An example of application in urban planning is described in (Chapman et al., 2018) - ten thousand Birmingham residents' posts on Twitter on the subject of urban green areas have been processed.

A personalized recommendation system based on user preferences (obtained by analyzing their reviews) and a physician's attributes (drug pricing preference, price, etc.) was presented in (Zhang et al., 2017).

Hate speech recognition in Twitter with high F_1 score of 0.93 has been achieved by deep learning (Badjatiya et al., 2017).

Forecasting stock price movements is very popular in the research community. In (Smailović et al., 2013) sentiment analysis of Twitter messages has shown that the polarity of the message may be a share price change indicator several days in advance.

This is just a few examples in a number of areas of sentiment analysis. Reviews of social events (such as festivals), marketing campaigns and company strategies are also often analyzed.

3 Challenges

Given the fact that the sentiment analysis is carried out over an unstructured text full of context, the success of the analysis depends on a number of factors. Some of the obstacles to successful analysis are the distinction between fact and opinion, syntax errors, negations, comparative sentences, irony and sarcasm recognition, fake opinions, context dependency, domain dependency, lack of resources for non-English languages, relatively small number of annotated datasets for supervised machine learning and the unsuccessfulness of the model in the domain where it was not developed.

Detecting subjectivity is not trivial to humans because it is a subjective task in itself - the sentence is for some neutral, and for someone not because of different levels of expertise on the topic or a different interpretation of the sentence (Chaturvedi et al., 2018).

Social network users are writing more freely and often do not care about writing proper words. Sometimes they deliberately write the words incorrectly to point out the thought (for example, they write the word *baaaad*, whose meanings are *very bad*) or they replace the letters by number (*2day* instead of *today*). Shortening is common (e.g. *whatews* instead of *whatever*), and it is not unusual for the two short words to join in a new one (e.g. *crunk* from *crazy / drunk*). In addition, the text can be filled with jargon and slang (Singh & Kumari, 2016).

Negations are an almost unavoidable challenge in the sentiment analysis because they most often change the polarity of opinion.

A special problem is irony and sarcasm that even people can not always understand, and this is a big problem for the computer during the analysis (Maynard & Greenwood, 2014).

The increase in the number of fake opinions is also one of the major problems (Jindal & Liu, 2008). Researchers try to find ways of detecting such reviews as they can significantly affect the outcome. One of them is the analysis of unusually positive interactions of potential spammers with ordinary users in order to make their opinion more influential (Choo, Yu & Chi, 2017).

Depending on the context, one and the same word can be both positive and negative and also neutral. For example, word *long* is basically neutral (*long hair*) but becomes negative in the example of *long queue* and positive in *long battery*. The word context is often far-off the word it depends on in the sentence (Chaturvedi et al., 2018), so it is difficult to understand when looking at the phrases in which words stand next to each other.

Due to hard and time consuming work, a relatively small number of labeled data sets required for supervised machine learning are evident (Mohammad, 2016).

Most of the research in this area is focused on texts written in English, and consequently the largest number of resources (labeled data sets, dictionaries ...) is developed for that language (Madhoushi, Hamdan & Zainudin, 2015). To overcome this issue, researchers perform automatic translations in which part of the text semantics is lost and the success of the analysis is reduced.

Sentiment analysis often depends on the problem domain, so a successful developed model in one domain may be unsuccessful when applied to another. The four major challenges of applying a model to a second domain are the scarcity of words (the target domain has words that do not exist or are rare in the original domain), polysemy (the same word has different meanings in two domains), divergence of features (unmatching in the original and target domains leads to classifier's misclassification) and divergence of polarity (word *easy* can be positive in one and negative in the other domain) (Al-Moslmi et al., 2017).

4 Methods and Approaches

For a successful sentiment analysis, proper preprocessing and feature engineering is required. Then the application of the techniques and methods of one of the two fundamental approaches (lexicon-based and machine learning) to the sentiment analysis follows (Figure 1). The lexicon-based approach has two variants: using a dictionary or a corpus-based approach. Depending on whether there are annotated learning data sets or not, machine learning can be supervised or unsupervised.

4.1 Data Preprocessing and Feature Engineering

Data may contain unwanted elements that slow down the learning process and reduce the efficiency of the developed model. Therefore, data is subject to preprocessing that partly depends on the data source. For example, preprocessing messages from Twitter may involve removing URLs (links to web sites do not contain useful information, and their removal can reduce the size of features), removing usernames (starting with @), removing tag #, removing duplicates and replies (containing the RT tag), compressing the 'extended' words (consecutive repeating of the characters, which increases the intensity of sentiment), converting capital letters to lowercase, and word stemming.

For the majority of machine learning applications, the main task in classifying sentiment is to select or extract features (Liu, 2010). In the sentiment analysis, feature engineering, among other things, uses (Pang & Lee, 2008):

• Terms and their frequencies - the term can be a single word or phrase; their appearance in the text is counted,

• Part-of-speech – tagging words by types (adjective, adverb, noun ...) can represent a rough form of clarifying the meaning of the word,

• Words and phrases of opinion – opinion is expressed by them,

• Negations - usually change the polarity of opinion.

By choosing features, one tries to reduce the dimensionality of data. Unlike this approach, feature transformation methods construct a smaller set of features based on the original set (Medhat, Hassan & Korashy, 2014). Because of the synonymy (different words of the similar or the same meaning) and polysemy (one word with more meanings) choosing features can be a slightly worse solution because the construction of new features reduces these two problems.

4.2 Lexicon-Based Approach

The lexicon-based approach refers to ways of creating a list of words (and phrases) that express sentiment. Although the word list can be made manually, in practice it is not applied alone because the process is both strenuous and time-consuming. Manual approach

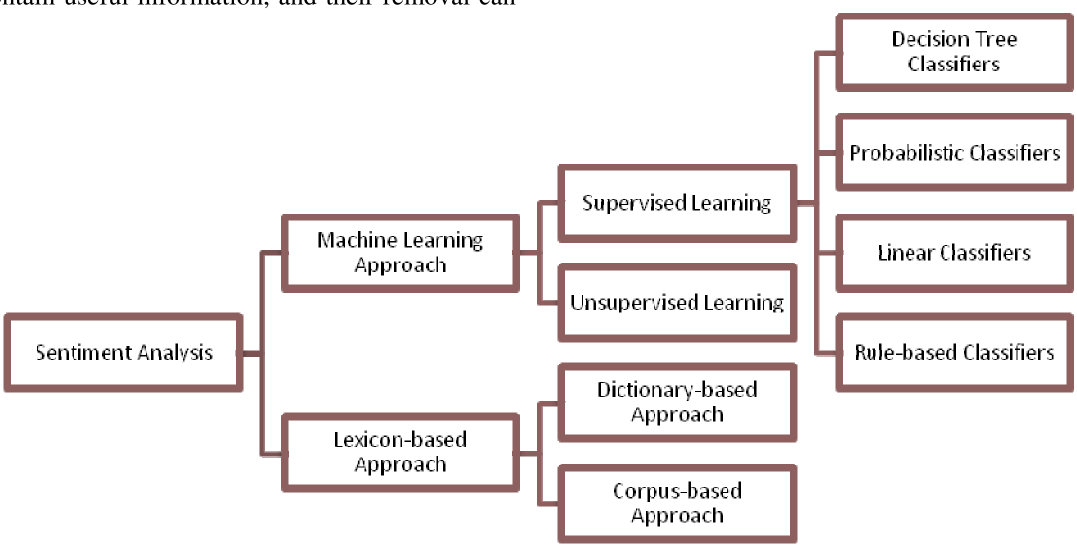

Figure 1. Approaches to sentiment analysis (adapted from (Medhat, Hassan & Korashy, 2014))

is commonly used to check automated approach since they create errors (Liu, 2015, p. 189). The dictionary-based approach uses a small set of manually collected words of known sentiment orientation that is then expanded using the available synonyms and antonyms dictionary (such as WordNet) to create a word list. The procedure is repeated with new words and ends when no new word can be found.

The corpus-based approach also begins with a list of words of known orientation and can be expanded with syntactic patterns (for example, two adjectives connected by conjunction *and* most likely are of the same orientation). The starting word list can also be generated by statistical methods. (Medhat, Hassan & Korashy, 2014)

In order to improve results, lexicon-based approach can be combined with machine learning methods. For example, in (Appel et al., 2018), a hybrid approach is used – it uses sentiment dictionary, semantic rules, language variables, negation and ambiguity handling.

4.3 Machine Learning

Appropriate learning examples are necessary in supervised machine learning. Namely, classifiers learn from examples, and then what he learned applies on new, test data. Usually, classification algorithms work on the following principle: algorithm learns from features - if a feature consistently falls into the same category of data classification while training, then it is a good indicator of that category and will be used in the testing phase.

Decision-making trees, linear classifiers (such as logistic regression), rule-based classifiers, and probabilistic classifiers (such as Naive Bayes or Maximum Entropy) can be used in sentiment analysis. The most commonly used are the Support Vector Machines and the naive Bayesian classifier (Ravi & Ravi, 2015). The Support Vector Machines because of its very good performance is often a basic method for comparing with new approaches.

Unsupervised machine learning tries to find structure in unlabeled data. Pointwise mutual information, latent semantic analysis (LSA), and Hidden Markov Models are some of the methods of unsupervised learning.

4.4 Dataless Classification

Dataless classification deals with a problem of lack of labeled learning data sets. It starts from the assumption that even the categories themselves carry a lot of information. People can easily decide whether sentence *Luka Modrić has been named as the best footballer in the world* belongs to the category of sport or politics. Dataless classification is a learning protocol that uses general knowledge to train the classifier without the need for any labeled data. Like humans, the data classifier interprets a sequence of words as a set of semantic concepts.

Using Wikipedia as a source of general knowledge and applying the Nearest Neighbors method, 85.29% accuracy was obtained on a set of 20 Newsgroups and 88.62% accuracy on Yahoo! Answers (Chang et al., 2008).

In (Song & Roth, 2014) the bootstrap dataless classification is performed with results comparable to the supervised classification on thousands of labeled examples.

In order to circumvent the large knowledge database (such as Wikipedia) learning problem, a new approach is introduced in (Chen et al., 2015). It uses a descriptive LDA (Latent Dirichlet Allocation) method that performs a dataless classification using only words that describe categories and undeclared documents. The LDA was used to obtain 30 latent themes based on the target data set, after which the category tag was manually assigned to each latent topic and were eventually manually tagged by category tags in each latent topic of the descriptor word. Authors state that the descriptive word list could be obtained in other ways, for example, from the dictionary.

4.5 Deep Learning

Deep learning has appeared a decade ago and has become a powerful machine learning technique with very good results in many areas of application, such as computer vision and natural text processing. Deep learning has recently become increasingly popular in sentiment analysis. Deep learning uses artificial neural networks in learning tasks through the network of multiple layers (Zhang, Wang & Liu, 2018).

Numerous deep learning models use word embedding for their inputs. The word embedding is a technique that converts words into vectors of real numbers in the vector space where words of similar meanings are closer to each other (Paltoglou & Thelwall, 2013). New words that were not in the training phase can be properly classified by similar words, which is an advantage over the approach that uses bag of words (Goldberg, 2016). Two commonly used word embedding systems are Global Vector (GloVe) - an unsupervised learning algorithm in which training is performed over a global word matrix (Pennington, Socher & Manning, 2014) and Word2Vec - which uses two models - one for predicting targeting word based on context (continuous Bag-of-Words model) and the other (Skip-Gram model) vice versa: based on the target word model it predicts the context (Mikolov et al., 2013).

5 Emotion Mining

Researches aimed at classifying text by polarity (positive / negative) dominate in sentiment analysis. However, more specific approach can be performed: according to emotions. This approach is only in the early phase of the research community (Plaza-del-Arco, 2018).

The initial problem in such an analysis is what emotions exist, and that will determine the classification classes. Among the many emotional models developed by psychologists, two of these are popular in computer recognition of emotions. The first is Eckman's model that differentiates the six basic emotions: anger, disgust, fear, happiness, sadness and surprise. The other is Plutchik's wheel of emotions with eight basic emotions: anger, anticipation, joy, trust, fear, surprise, sadness and disgust. Two emotions are on opposite sides (e.g. joy vs. sadness). The basic emotions can be combined to create new ones, and the color shows their intensity (Figure 2).

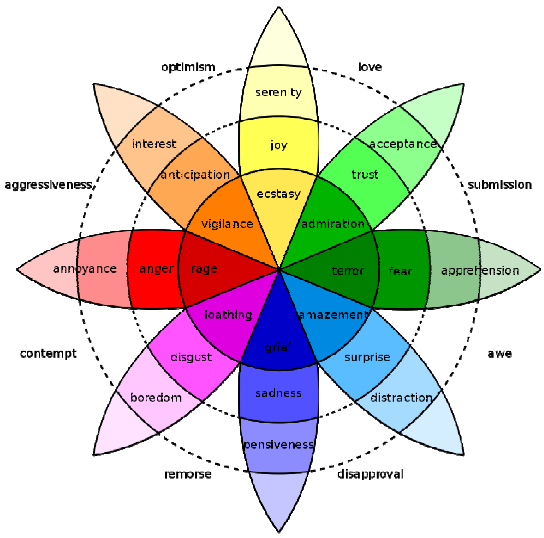

Figure 2. Plutchik's wheel of emotions (Wikimedia Commons)

In (Bravo-Marquez, Mendoza & Poblete, 2013) it has been shown that using the NRC emotion lexicon (based on the Plutchik's model) when combined with other lexicons that contain numerically expressed polarizations, they get better results in classification of two data sets (Stanford Twitter Sentiment and Sanders) that contain Twitter messages.

In (Agrawal & Papagelis, 2018), the goal was to overcome one lack of the learning method of word representation. Namely, most methods are based on the distribution hypothesis in linguistics, according to which the used words that appear in the same context tend to have similar meanings. As a result, the appearance of emotionally different words, such as *happy* and *sad* in similar contexts would have a greater similarity in the vector space than emotionally similar words, such as *happy* and *joy*. WordNetAffect and NRC dictionaries were used to improve word representation in the vector space, and the results surpassed several general purpose word representations (GloVe and Word2Vec).

A model based on emotional state data presented in (Martins et al., 2018) can be used to predict errors or to recommend people to stop working in certain states for immediate work to avoid dangerous mistakes. The model is created by combining text typing log (includes typing time data, total and average number of errors, time between typing, repeating frequency of typing the same letters, etc.) with the analysis of emotions made over the text itself.

Lexical categories such as frequency of use of personal pronouns and verbal past tenses correlate with personality traits and can be used to detect them. In (Mohammad & Kiritchenko, 2013) it has been shown that emotions are significant personality indicators and can serve as a means of detection. In this research, the NRC hashtags emotion dictionary, Osgood dimension dictionary and vocabulary of specificities were used Classification was performed by Supporting Vector Machines. The use of fine emotional features has led to a statistically significant improvement over the comparable basic model.

Lack of labeled data is one of the important brakes for faster development of emotion mining. For polarity analysis there is a much larger amount of labeled data, primarily due to the fact that this area is being investigated longer, but also because of the less effort in tagging the text because one of the three categories should be considered, while in emotional tagging it can be expected more than the default six or eight (or some other number, depending on the selected model) emotions. One sentence can simultaneously contain both surprise and anger, and perhaps fear, so labelling consumes a lot more time. On the other hand, the successful application of sentiment analysis (several selected examples are briefly described in this chapter) could stimulate more interest in research in this direction.

6 Conclusion

By popularizing Web 2.0 technologies, the internet has become a rich source of user-generated content. Users from all over the world express and publicly share their opinions on different topics. Manual analysis of large amounts of such data is impossible, so a reasonable need for their computer processing has emerged. Sentiment analysis collects and processes people's opinions and attitudes toward products, services, politics, social events, marketing campaigns and company strategies. Reviews (from sources such as TripAdvisor, Amazon and IMDB) and social network posts (mostly from Twitter and Facebook) are categories of textual documents that are the most interesting for sentiment analysis.

In the sentiment analysis, unstructured text is processed, and it brings numerous issues in computer processing. Wrongly written words, irony and sarcasm, complex dependent sentences, comparative sentences, recognition of opinions in relation to neutral (objective) sentences make it difficult for a successful analysis. Furthermore, there is an evident lack of human-labeled data sets required by machine learning algorithms. This problem is more prominent beyond the English speaking area, so researchers are performing automatic translations in which part of the text semantics is lost and the success of the analysis is reduced. An additional problem is that models that successfully classify text in one problem domain show significantly worse results in the other.

In order to overcome these problems, different methods and techniques that are based on machine learning or lexical approach are used. In a recent time a deep learning method is becoming more and more popular in sentiment analysis. Deep learning applies artificial neural networks that recognize semantic relationships in the text and give results comparable to methods that are longer present in machine learning.

Although deep learning is a relatively new approach in sentiment analysis, it is already promising in polarity analysis. However, as deep learning needs larger sets of labeled learning examples, it may have less success in emotion mining.

A dataless classification that uses general knowledge to train a classifier seems like an interesting approach in situations where there are no labeled learning data sets.

Compared to isolated use of a single method, approaches that combine multiple methods and techniques (for example the use of vocabulary and machine learning) can be a good way to overcome the challenges of sentiment analysis.

Sentiment analysis is an area that is intensifying, thanks to the wealth of data sources, but also to many useful applications. The demands placed in front of it are getting higher and higher. At first it was enough to recognize whether the document is predominantly positive or negative. Over a time a need to get polarity at the level of an aspect of a product or service arose, which is significantly more difficult. Likewise, instead of recognizing positivity or negativity, researchers are increasingly investigating recognition of concrete emotions in the text, and in the future it is likely that new challenges and tasks will be placed on the sentiment analysis.

Future research will focus on techniques and methods of feature selection and extraction, which are very obscurely described in this paper.

References

Agrawal, A., An, A. & Papagelis (2018). M. Learning Emotion-enriched Word Representations. In *Proceedings of the 27th International Conference on Computational Linguistics*, 950-961.

Al-Moslmi, T., Omar, N., Abdullah, S., & Albared, M. (2017). Approaches to cross-domain sentiment analysis: A systematic literature review. *IEEE Access*, 5, 16173-16192.

Appel, O. et al. (2018). Successes and challenges in developing a hybrid approach to sentiment analysis. *Applied Intelligence*, 48(5), 1176-1188.

Badjatiya, P. et al. (2017). Deep learning for hate speech detection in tweets. In *Proceedings of the 26th International Conference on World Wide Web Companion*, 759-760, International World Wide Web Conferences Steering Committee.

Bravo-Marquez, F., Mendoza, M. & Poblete, B. (2013). Combining strengths, emotions and polarities for boosting Twitter sentiment analysis. In *Proceedings of the Second International Workshop on Issues of Sentiment Discovery and Opinion Mining* (p. 2). ACM.

Chang, M.W. et al. (2008). Importance of Semantic Representation: Dataless Classification. In *Proceedings of the 23rd AAAI conference on Artificial intelligence*, Vol. 2, 830-835.

Chapman, L. et al. (2018). Investigating the emotional responses of individuals to urban green space using twitter data: A critical comparison of three different methods of sentiment analysis. *Urban Planning*, 3(1), 21-33.

Chaturvedi, I. et al. (2018). Distinguishing between facts and opinions for sentiment analysis: Survey and challenges. *Information Fusion*, 44, 65-77.

Chen, X. et al. (2015). Dataless text classification with descriptive LDA. In *Proceedings of the 29th AAAI conference on Artificial intelligence*, 2224-2231.

Choo, E., Yu, T. & Chi, M. (2017). Detecting opinion spammer groups and spam targets through community discovery and sentiment analysis. *Journal of Computer Security*, 25.3: 283-318.

Esuli, A. & Sebastiani, F. (2006). SENTIWORDNET: A Publicly Available Lexical Resource for Opinion Mining. In *5th Conference on Language Resources and Evaluation*, 417-422.

Goldberg, Y. (2016). A primer on neural network models for natural language processing. *Journal of Artificial Intelligence Research*, 57, 345-420.

Jindal, N. & Liu, B. (2008). Opinion spam and analysis. In *Proceedings of the 2008 international conference on web search and data mining*, ACM, 219-230.

Liu, B. (2010). Sentiment Analysis and Subjectivity. *Handbook of natural language processing*, 2, 627-666.

Liu, B. (2015). *Sentiment analysis: Mining opinions, sentiments, and emotions*. Cambridge University Press, New York.

Madhoushi, Z., Hamdan, A.R. & Zainudin, S. (2015). Sentiment analysis techniques in recent works. In *Science and Information Conference (SAI)*, IEEE.

Martins, R. et al. (2018). Predicting Performance Problems Through Emotional Analysis (Short Paper). In *7th Symposium on Languages, Applications and Technologies (SLATE 2018)*, Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.

Maynard, D. & Greenwood, M.A. (2014). Who cares about Sarcastic Tweets? Investigating the Impact of Sarcasm on Sentiment Analysis. *Lrec*, 4238-4243.

Medhat, W., Hassan, A. & Korashy, H. (2014). Sentiment analysis algorithms and applications: A survey. *Ain Shams Engineering Journal*, 5.4: 1093-1113.

Mikolov, T. et al. (2013). Efficient estimation of word representations in vector space. *arXiv preprint* arXiv:1301.3781.

Mohammad, S. M. (2016). Sentiment analysis: Detecting valence, emotions, and other affectual states from text. *Emotion measurement*, 201-237.

Mohammad, S. M. & Kiritchenko, S. (2013). Using nuances of emotion to identify personality. *Proceedings of ICWSM*, 27-30.

Paltoglou, G. & Thelwall, M. (2013). More than bag-of-words: Sentence-based document representation for sentiment analysis. In *Proceedings of the International Conference Recent Advances in Natural Language Processing RANLP 2013*, 546-552.

Pang, B. & Lee, L. (2008). Opinion mining and sentiment analysis. *Foundations and Trends® in Information Retrieval*, 2(1–2): 1-135.

Pennington, J., Socher, R. & Manning, C. (2014). Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 1532-1543.

Plaza-del-Arco, F. M. et al. (2018). Lexicon Adaptation for Spanish Emotion Mining. *Procesamiento del Lenguaje Natural*, 61, 117-124.

Ravi, K. & Ravi, V. (2015). A survey on opinion mining and sentiment analysis: tasks, approaches and applications. *Knowledge-Based Systems*, 89, 14-46.

Singh, T. & Kumari, M. (2016). Role of text pre-processing in twitter sentiment analysis. *Procedia Computer Science*, 89: 549-554.

Smailović, J. et al. (2013). Predictive sentiment analysis of tweets: A stock market application. In Human-computer interaction and knowledge discovery in complex, unstructured, *Big Data*, 77-88, Springer, Berlin, Heidelberg.

Song, Y. & Roth, D. (2014). On Dataless Hierarchical Text Classification. In *Proceedings of the 28th AAAI conference on Artificial intelligence*, Vol. 7, 1579-1585.

Vadicamo, L. et al. (2017). Cross-Media Learning for Image Sentiment Analysis in the Wild. In *ICCV Workshops*, 308-317.

Xu, C. et al. (2014). Visual sentiment prediction with deep convolutional neural networks. *arXiv preprint*, arXiv:1411.5731.

Zhang, L., Wang, S. & Liu, B. (2018). Deep learning for sentiment analysis: A survey. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 8(4), e1253.

Zhang, Y. et al. (2017). iDoctor: Personalized and professionalized medical recommendations based on hybrid matrix factorization. *Future Generation Computer Systems*, 66, 30-35.